The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Complexity of Verification in Incomplete Argumentation Frameworks

Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, Hilmar Schadrack
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany

Abstract

Abstract argumentation frameworks are a well-established formalism to model nonmonotonic reasoning processes. However, the standard model cannot express incomplete or conflicting knowledge about the state of a given argumentation. Previously, argumentation frameworks were extended to allow uncertainty regarding the set of attacks or the set of arguments. We combine both models into a model of general incompleteness, complement previous results on the complexity of the verification problem in incomplete argumentation frameworks, and provide a full complexity map covering all three models and all classical semantics. Our main result shows that the complexity of verifying the preferred semantics rises from coNP- to Σ_2^p-completeness when allowing uncertainty about either attacks or arguments, or both.

1 Introduction

Within the field of artificial intelligence, abstract argumentation frameworks have emerged as a useful methodology to represent and evaluate nonmonotonic logics. They allow to create a simple, directed graph from a defeasible knowledge base that consists of only arguments (nodes) and attacks (directed edges), then to identify sets of "acceptable" arguments in that graph, and finally to interpret these arguments' conclusions as models in the knowledge base. In this framework, when evaluating which arguments are acceptable in the graph, the internal structure of arguments is neglected, which accounts for the simplicity of the formalism.

Since Dung (1995) introduced his seminal model, many model extensions of argumentation frameworks have been proposed that allow to capture a wider and more fine-grained range of applications. This paper continues a line of research aimed at expressing unquantified uncertainty in an argumentation framework. Such *qualitative* uncertainty about the state of an argumentation framework was introduced by Coste-Marquis et al. (2007) and further studied by Baumeister, Neugebauer, and Rothe (2015) for the set of attacks and by Baumeister, Rothe, and Schadrack (2015) for the set of arguments. This paper is the first to allow uncertainty about both arguments and attacks simultaneously. Our main question is to examine how the complexity of verifying certain semantics (expressing which subsets of the arguments are acceptable in various ways) changes when asking whether they are satisfied *possibly* (in some completion of the incomplete graph) or *necessarily* (in all its completions).

A large body of previous work in abstract argumentation addresses *quantitative* uncertainty about the state of a given argumentation by using probabilities. Fuzzy argumentation frameworks (Janssen, Cock, and Vermeir 2008) replace the attack relation with a fuzzy relation, where each individual attack has a fuzzy value in $[0, 1]$ that represents the degree to which this attack holds. In a fuzzy argumentation framework, for two sets of arguments, the degree to which they attack each other can be determined. In probabilistic argumentation frameworks, Li, Oren, and Norman (2011) assume that a probability distribution over both arguments and attacks is given. Other approaches associate a probability with each set of arguments (Dung and Thang 2010; Rienstra 2012) to indicate whether all and only these arguments are active, or with each spanning subtree of the argument graph (Hunter 2014) to indicate that all and only the attacks contained in that subtree are active. In all these models, an interesting question is to determine the probability for a set of arguments to be acceptable. A different branch of research on probabilistic argumentation uses probabilities to represent the epistemic state of arguments, attacks, or sets of arguments, i.e., the belief in those elements (in terms of acceptance). Although technically similar, this approach has a completely different purpose than ours, which is the representation of *structural* uncertainty.

Another field that raises similar questions is that of dynamic change of argumentation frameworks. Previous work has examined how adding or deleting a set of arguments can alter the set of acceptable sets of arguments (Cayrol, de Saint-Cyr, and Lagasquie-Schiex 2010; Boella, Kaci, and van der Torre 2009), the complexity of computing the acceptability of a single argument after changing the arguments or attacks (Liao, Jin, and Koons 2011), or enforcement of a set of arguments (Baumann and Brewka 2010; Wallner, Niskanen, and Järvisalo 2016; Coste-Marquis et al. 2015), where the question is how much a given argumentation framework needs to be modified to make the given set of arguments acceptable.

In the remainder of this paper, we give the required background in abstract argumentation (Section 2), introduce incomplete argumentation frameworks as a generalization of attack- and argument-incomplete argumentation frameworks (Section 3), followed by a full complexity analysis of verification in all three models (Section 4), and we discuss our results and possible tasks for future work (Section 5).

2 Preliminaries

We start by defining argumentation frameworks due to Dung (1995), mostly following the notation by Dunne and Wooldridge (2009).

Definition 1. An *argumentation framework* AF is a pair $\langle \mathcal{A}, \mathcal{R} \rangle$ consisting of a finite set $\mathcal{A}$ of *arguments* and a binary *attack relation* $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ on the arguments. We say that a *attacks* b if $(a, b) \in \mathcal{R}$.

An argumentation framework $AF = \langle \mathcal{A}, \mathcal{R} \rangle$ can be displayed as a directed graph (V, E) by identifying arguments with vertices and attacks with directed edges: $V = \mathcal{A}$ and $E = \mathcal{R}$.

Example 2. Consider an argumentation framework $AF = \langle \mathcal{A}, \mathcal{R} \rangle$ with $\mathcal{A} = \{a, b, c, d, e\}$ and $\mathcal{R} = \{(a, b), (b, a), (b, c), (c, d), (d, e), (e, c), (e, d), (e, e)\}$. Its graph representation is given in Figure 1.

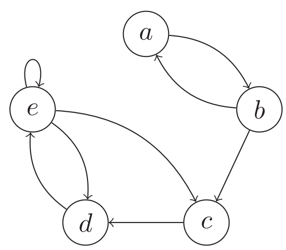

Figure 1: Argumentation framework for Example 2

The main objective in abstract argumentation is to identify sets of arguments that are simultaneously acceptable. Various *semantics* were defined in the literature that impose different acceptability conditions for sets of arguments. We cover all semantics that were defined in the seminal paper by Dung (1995). They are formalized in Definition 3, after introducing some necessary notions.

An argument $a \in \mathcal{A}$ is *defended by* $S \subseteq \mathcal{A}$ if, for each $b \in \mathcal{A}$ with $(b, a) \in \mathcal{R}$, there is a $c \in S$ such that $(c, b) \in \mathcal{R}$. For an argumentation framework AF, the *characteristic function* $F_{AF} : 2^{\mathcal{A}} \to 2^{\mathcal{A}}$ maps each set S of arguments to the set of arguments that are defended by S, i.e., $F_{AF}(S) = \{a \in \mathcal{A} \mid a \text{ is defended by } S\}$. The characteristic function always has a least fixed point, since it is monotonic with respect to set inclusion. Let F^k_{AF} denote the k-fold composition of F_{AF}, and let F^*_{AF} denote the infinite composition, which yields the fixed points of F_{AF}.

Definition 3. Let $AF = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework. A set $S \subseteq \mathcal{A}$ is

- *conflict-free* if $(a, b) \notin \mathcal{R}$ for all $a, b \in S$.

A conflict-free set $S \subseteq \mathcal{A}$ is

- *admissible* if $S \subseteq F_{AF}(S)$,
- *complete* if $S = F_{AF}(S)$,
- *grounded* if $S = F^*_{AF}(\emptyset)$, i.e., S is the least fixed point of F_{AF},
- *preferred* if $S \subseteq F_{AF}(S)$ and there is no admissible set $S' \supset S$, and
- *stable* if for every $b \in \mathcal{A} \setminus S$ there is an $a \in S$ with $(a, b) \in \mathcal{R}$.

Among these properties, conflict-freeness and admissibility are typically considered to be basic requirements while the others are "real" semantics—for the sake of convenience, however, we will not always distinguish between basic properties and semantics.

It is obvious that the grounded set is unique and complete and that every complete set is admissible. The work of Dung (1995) further provides that there always is a conflict-free, admissible, complete, grounded, and preferred set, but there may be no stable set. Also, every stable set is preferred, every preferred set is complete, and every admissible set is conflict-free.

We assume the reader to be familiar with the complexity classes of the polynomial hierarchy, in particular, P, NP, coNP, and $\Sigma^p_2 = \mathrm{NP}^{\mathrm{NP}}$, as well as the concepts of hardness and completeness. For an introduction, see, e.g., the books by Papadimitriou (1995) and Rothe (2005).

Dunne and Wooldridge (2009) defined decision problems regarding the existence or status of acceptable arguments. We focus on the verification problem, which is parameterized by one of the semantics (denoted $\mathbf{s}$) defined above and asks whether a given set of arguments is an *extension* of the argumentation framework with respect to that semantics, i.e., whether it satisfies the conditions imposed by that semantics. As shorthands, we may use CF for *conflict-free*, AD for *admissible*, CP for *complete*, GR for *grounded*, PR for *preferred*, and ST for *stable* semantics.

$\mathbf{s}$-VERIFICATION ($\mathbf{s}$-VER)	
Given:	An argumentation framework $\langle \mathcal{A}, \mathcal{R} \rangle$ and a subset $S \subseteq \mathcal{A}$.
Question:	Is S an $\mathbf{s}$ extension of AF?

The problem PR-VERIFICATION was shown to be coNP-complete by Dimopoulos and Torres (1996), but Dung (1995) established polynomial-time algorithms for verifying the other semantics from Definition 3.

3 Incomplete Argumentation Frameworks

In our model of incomplete argumentation framework, both the set of arguments and the set of attacks are split into a *definite* and a *possible* set, which represent the elements that are known to exist, respectively, which may or may not exist.

Definition 4. An *incomplete argumentation framework* is a quadruple $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, where $\mathcal{A}$ and $\mathcal{A}^?$ are disjoint sets of arguments and $\mathcal{R}$ and $\mathcal{R}^?$ are disjoint subsets of $(\mathcal{A} \cup \mathcal{A}^?) \times (\mathcal{A} \cup \mathcal{A}^?)$. $\mathcal{A}$ is the set of arguments that are known to definitely exist, while $\mathcal{A}^?$ contains all possible additional arguments not (yet) known to exist. Similarly, $\mathcal{R}$ is the set of attacks that are known to definitely exist (as long as both incident arguments turn out to exist), while $\mathcal{R}^?$ contains all possible additional attacks not (yet) known to exist.

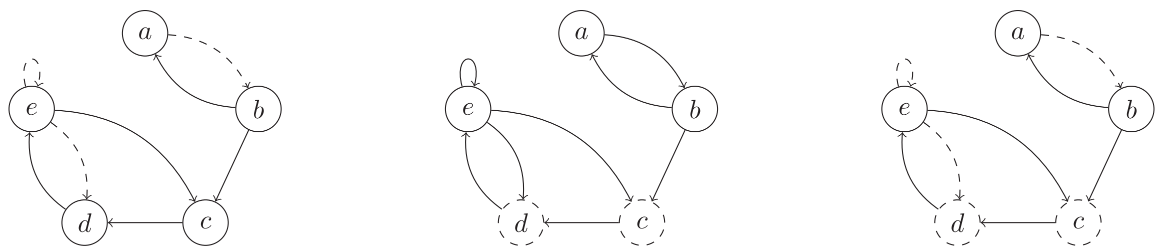

(a) An attack-incomplete AF (b) An argument-incomplete AF (c) An incomplete AF

Figure 2: Some examples of incomplete argumentation frameworks

Example 5. Figure 2 displays graph representations of three incomplete argumentation frameworks, where definite elements are displayed as usual and possible elements are displayed as dashed circles or arcs. Elements that are known to not exist are not displayed.

The incomplete argumentation framework in Figure 2a has no uncertainty regarding the arguments, while the one in Figure 2b has no uncertainty regarding the attacks. The incomplete argumentation framework in Figure 2c combines the uncertainty of the other two.

An incomplete argumentation framework $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ can be seen as a representation of a finite universe of possible worlds, where each world corresponds to a single argumentation framework (without uncertainty), in which each possible argument in $\mathcal{A}^?$ and each possible attack in $\mathcal{R}^?$ is either included or excluded. Such an argumentation framework is called a *completion of* $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$. When excluding a possible argument, all its incident attacks are also automatically excluded: For a set $\mathcal{A}^*$ of arguments with $\mathcal{A} \subseteq \mathcal{A}^* \subseteq \mathcal{A} \cup \mathcal{A}^?$, the *restriction of a relation* $\mathcal{R}$ *to* $\mathcal{A}^*$ is $\mathcal{R}|_{\mathcal{A}^*} = \{(a,b) \in \mathcal{R} \mid a,b \in \mathcal{A}^*\}$.

Definition 6. Let $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete argumentation framework. An argumentation framework $IAF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ with $\mathcal{A} \subseteq \mathcal{A}^* \subseteq \mathcal{A} \cup \mathcal{A}^?$ and $\mathcal{R}|_{\mathcal{A}^*} \subseteq \mathcal{R}^* \subseteq \left(\mathcal{R} \cup \mathcal{R}^?\right)|_{\mathcal{A}^*}$ is called a *completion of* IAF.

In general, the number of possible completions is exponential in the size of the incomplete argumentation framework—it is at most $2^{|\mathcal{R}^?|+|\mathcal{A}^?|}$, but may be slightly lower: Since excluding possible arguments may implicitly also exclude possible attacks, it may be that some of the completions coincide.

Example 7. Continuing Example 5, the incomplete argumentation frameworks in Figures 2a and 2b have $2^3 = 8$ and $2^2 = 4$ completions, respectively. The incomplete argumentation framework in Figure 2c has 24 completions: $2^4 = 16$ that include argument d, and another $2^3 = 8$ that exclude d, since, in the latter case, the attack (e,d) is not available.

In an incomplete argumentation framework IAF, we say that a property defined for standard argumentation frameworks (e.g., a semantics) holds *possibly* if there exists a completion IAF^* of IAF for which the property holds, and a property holds *necessarily* if it holds for all completions of IAF. Thus we can define two variants of the verification problem in the incomplete case for each given semantics $\mathbf{s}$:[1]

$\mathbf{s}$-INC-POSSIBLE-VERIFICATION ($\mathbf{s}$-INCPV)	
Given:	An incomplete argumentation framework $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and a set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$.
Question:	Is there a completion $IAF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ of IAF such that $S|_{\mathcal{A}^*} = S \cap \mathcal{A}^*$ is an $\mathbf{s}$ extension of IAF^*?

$\mathbf{s}$-INC-NECESSARY-VERIFICATION ($\mathbf{s}$-INCNV)	
Given:	An incomplete argumentation framework $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and a set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$.
Question:	For all completions $IAF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ of IAF, is $S|_{\mathcal{A}^*} = S \cap \mathcal{A}^*$ an $\mathbf{s}$ extension of IAF^*?

Both problems are potentially harder than standard verification, since they add an existential (respectively, universal) quantifier over a potentially exponential space of solutions.

Incomplete argumentation frameworks are a generalization of both pure models of incomplete argumentation frameworks. Fixing $\mathcal{A}^? = \emptyset$ in Definitions 4 and 6 yields exactly the class of attack-incomplete argumentation frameworks as proposed by Coste-Marquis et al. (2007) (and further studied by Baumeister, Neugebauer, and Rothe (2015)), and fixing $\mathcal{R}^? = \emptyset$ yields exactly the class of argument-incomplete argumentation frameworks as proposed by Baumeister, Rothe, and Schadrack (2015). In attack-incomplete argumentation frameworks, the set of possible arguments can be omitted and it can be written as $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$. Likewise, $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$ denotes a purely argument-incomplete argumentation framework. Also, there are distinct possible and necessary variants of the verification problem for both pure models of incompleteness, which were introduced by Baumeister, Neugebauer, and Rothe (2015) and Baumeister, Rothe, and Schadrack (2015). In the attack-incomplete model, we write $\mathbf{s}$-ATTINCPV and $\mathbf{s}$-ATTINCNV for possible and necessary verification, respectively, and $\mathbf{s}$-ARGINCPV and $\mathbf{s}$-ARGINCNV in the argument-incomplete model.

[1]Maher (2016) studies resistance to corruption in strategic argumentation. While instances in his model and in our argument-incomplete argumentation frameworks are technically similar, his results do not carry over to our problems. One difference is that he focuses on credulous or skeptical acceptance of specific arguments, whereas we consider verification of entire extensions.

4 Complexity of Verification

In this section, we provide a full complexity analysis of the verification problem in the general model of incomplete argumentation framework and also fill the gaps in complexity for both pure models. Since the combined model is a generalization of both pure models, all upper bounds in the general incompleteness model straightforwardly provide the same upper bound for both pure incompleteness models. Likewise, all lower bounds in either of the pure incompleteness models immediately give the same lower bound for the general model. We start by providing some simple upper bounds for the complexity in the general incompleteness model.

Theorem 8. *For* $\mathrm{s} \in \{\textsc{ad}, \textsc{st}, \textsc{cp}, \textsc{gr}\}$, s-INCPV *is in* NP. *Moreover,* PR-INCPV *is in* Σ_2^p *and* PR-INCNV *is in* coNP.

Proof. The results follow directly from the quantifier representations of the given problems: We start with an existential quantifier for possible verification, and with a universal quantifier for necessary verification. For $\mathrm{s} \in \{\textsc{ad}, \textsc{st}, \textsc{cp}, \textsc{gr}\}$, it can be checked in polynomial time whether the given subset is an s extension. Moreover, the standard verification problem for the preferred semantics belongs to coNP; hence, it can be written as a universal quantifier followed by a statement checkable in polynomial time. Therefore, this polynomial-time predicate is preceded first by an existential quantifier (guessing a completion) and then a universal quantifier (verifying preferredness) in the case of PR-INCPV, and it is preceded by two universal quantifiers that can be collapsed to one such quantifier for PR-INCNV. □

Baumeister, Rothe, and Schadrack (2015) proved that s-ARGINCPV is NP-hard for $\mathrm{s} \in \{\textsc{ad}, \textsc{st}, \textsc{cp}, \textsc{gr}\}$, and the work by Dimopoulos and Torres (1996) yields that PR-VERIFICATION is coNP-hard. These hardness results carry over to the general model and coincide with the respective upper bounds from Theorem 8.

Corollary 9. *For* $\mathrm{s} \in \{\textsc{ad}, \textsc{st}, \textsc{cp}, \textsc{gr}\}$, s-INCPV *is* NP-*complete. Moreover,* PR-INCNV *is* coNP-*complete.*

4.1 Upper Bounds

In this section, we provide all remaining upper bounds on the complexity of possible and necessary verification. We start by proving that verifying conflict-freeness remains easy in the general model.

Theorem 10. CF-INCPV *and* CF-INCNV *both are in* P.

Proof. Given an incomplete argumentation framework $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and a set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ of arguments, S is possibly conflict-free in IAF if and only if $S|_{\mathcal{A}}$ is conflict-free in the *minimal completion* $\langle \mathcal{A}, \mathcal{R}|_{\mathcal{A}} \rangle$ of IAF, which discards all possible arguments and attacks. Similarly, S is necessarily conflict-free in IAF if and only if S is conflict-free in the *maximal completion* $\langle \mathcal{A} \cup \mathcal{A}^?, \mathcal{R} \cup \mathcal{R}^? \rangle$ of IAF, which includes all possible arguments and attacks. Since both the minimal and the maximal completion can clearly be constructed in polynomial time, we have P membership for both problems. □

In the remainder of Section 4.1, we will, step by step, show that all upper bounds from the standard model without incompleteness are preserved by the necessary verification variant in all three incompleteness models and for all considered semantics.

Theorem 11. AD-ARGINCNV *and* ST-ARGINCNV *both are in* P.

Proof. Let $I = (\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle, S)$ be an instance of AD-ARGINCNV. If S is not necessarily conflict-free in $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$, it is not necessarily admissible in $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$, either. Since CF-ARGINCNV is in P (Baumeister, Rothe, and Schadrack 2015), this can be checked in polynomial time. In the following, we may assume that S is necessarily conflict-free.

Let $\mathcal{A}_0 = \mathcal{A} \cup (\mathcal{A}^? \setminus S)$ and $C_0 = \langle \mathcal{A}_0, \mathcal{R}|_{\mathcal{A}_0} \rangle$, and for each argument $a \in \mathcal{A}^? \cap S$, let $\mathcal{A}_a = \mathcal{A}_0 \cup \{a\}$ and $C_a = \langle \mathcal{A}_a, \mathcal{R}|_{\mathcal{A}_a} \rangle$. If, for some $x \in \{0\} \cup (\mathcal{A}^? \cap S)$, $S|_{\mathcal{A}_x}$ is not admissible in the completion C_x, we clearly have $I \notin$ AD-ARGINCNV. Since the number of these completions is bounded by the number of arguments (plus one), this can again be verified in polynomial time. We may now assume that, in each completion C_x, $S|_{\mathcal{A}_x}$ is admissible.

Note that each of these completions includes *all* possible attacks against the respective set $S|_{\mathcal{A}_x}$, because the completions include all possibly harmful arguments (members of $\mathcal{A}_0$) and because there cannot be any attacks among members of S. This yields that $S|_{\mathcal{A}_0}$ defends all attacks against its elements in *any* completion, and, for all $a \in \mathcal{A}^? \cap S$, $S|_{\mathcal{A}_a}$ defends all attacks against a in *any* completion. Finally, since in any completion $C^* = \langle \mathcal{A}^*, \mathcal{R}|_{\mathcal{A}^*} \rangle$, it holds that $S|_{\mathcal{A}^*} \subseteq \bigcup_x S|_{\mathcal{A}_x}$, we can conclude that each element of $S|_{\mathcal{A}^*}$ is defended by $S|_{\mathcal{A}^*}$ in C^*, so S is necessarily admissible in $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$ and $I \in$ AD-ARGINCNV.

ST-ARGINCNV $\in$ P can be proven with the same construction and an analogous argumentation. □

The previous result can be lifted to the general incompleteness model. We sketch the proof idea.

Theorem 12. AD-INCNV *and* ST-INCNV *both are in* P.

Proof (sketch). For an instance (IAF, S) of AD-INCNV, reduce the problem to AD-ARGINCNV using a single critical "pessimistic" argument-incomplete argumentation framework IAF_S^{pes}, which is obtained by including each and only those attacks that target S. P membership then follows from a straightforward proof of the equivalence $(IAF, S) \in$ AD-INCNV $\Leftrightarrow (IAF_S^{pes}, S) \in$ AD-ARGINCNV. An analogous proof can be used for the stable semantics. □

Turning to the complete and grounded semantics, we can successively prove P membership of CP-INCNV and GR-INCNV in Theorems 13 and 17, respectively.

Theorem 13. CP-INCNV *is in* P.

Proof. Let (IAF, S) with $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an instance of CP-INCNV. Since AD-INCNV $\in$ P, we may assume that S is necessarily admissible in IAF. We clearly have $(IAF, S) \notin$ CP-INCNV if and only if there is at least one argument outside of S that is defended by S in some completion of IAF. It remains to show how to check this.

If all arguments $a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$ are definitely attacked by S, i.e., $(b_a, a) \in \mathcal{R}$ for each such argument a and some corresponding $b_a \in S$, then S is necessarily stable and therefore necessarily complete, and we are done. Now assume this is not the case and let $a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$ be any argument outside of S that is not definitely attacked by S, i.e., $(b, a) \notin \mathcal{R}$ for all $b \in S \cap \mathcal{A}$ (if a were attacked by S, it clearly could not be defended by S in any completion). Let $Att(a) = \{b \in \mathcal{A} \cup \mathcal{A}^? \mid (b, a) \in \mathcal{R}\}$ be the set of all arguments with a definite attack against a. Further, let $\mathcal{R}_a = \mathcal{R} \cup \{(b, c) \in \mathcal{R}^? \mid b \in S \text{ and } c \in Att(a) \setminus \{a\}\}$ be the set of attacks that includes all and only those possible attacks for which the attacker is in S and the target is an attacker of a.

Consider now the completion $C_a = \langle \mathcal{A}_a, \mathcal{R}_a|_{\mathcal{A}_a} \rangle$ where $\mathcal{A}_a = \mathcal{A} \cup \{a\} \cup \{b \in \mathcal{A}^? \mid (b, a) \notin \mathcal{R}_a\}$, i.e., C_a uses the attack relation $\mathcal{R}_a$ and includes a and exactly those possible arguments that do not attack a (in $\mathcal{R}_a$). If, for any of these completions, a is defended by S in C_a, then S is not complete in C_a and therefore not necessarily complete. If, on the other hand, each argument a is not defended by S in the respective completion C_a, then none of these arguments are possibly defended by S, and therefore, S is necessarily complete: Assume that a is not defended by S in C_a, i.e., there is some $b \in \mathcal{A}_a$ with $(b, a) \in \mathcal{R}_a|_{\mathcal{A}_a}$ and S does not attack b in C_a. By construction of C_a, we know that b is a definite argument, i.e., $b \in \mathcal{A}$, and (b, a) is a definite attack, i.e., $(b, a) \in \mathcal{R}$, so b attacks a in any completion that contains a. Also, in all completions, S either does not defend a against b, or S attacks a, since all possible arguments in S either attack a or are already included in C_a. So, a is not possibly defended by S.

All steps taken can clearly be performed in polynomial time. This completes the proof. □

This result immediately yields P membership for the corresponding problem in the argument-incomplete model.

Corollary 14. CP-ARGINCNV *is in* P.

Next, we introduce the notion of ungrounded completion of an incomplete argumentation framework as a tool to prove P membership of GR-INCNV.

Definition 15. Let $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete argumentation framework and $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ be a set of arguments in IAF. The *ungrounded completion* IAF_S^{ungr} of IAF for S is obtained by the following algorithm.

1. Let $\mathcal{R}_0 = \mathcal{R} \cup \{(a, b) \in \mathcal{R}^? \mid b \in S\}$.
2. Let $G_0 = \emptyset$, $\mathcal{A}_0^? = \mathcal{A}^?$, $IAF_0 = \langle \mathcal{A}, \mathcal{A}_0^?, \mathcal{R}_0 \rangle$, and $i = 0$.
3. Let Max_i be the maximal completion of IAF_i and let $X_i \subseteq S$ be the set of arguments in S that are defended by G_i in Max_i, i.e., $X_i = F_{\text{Max}_i}(G_i) \cap S$. Add the definite arguments in X_i to G_i and exclude the possible arguments in X_i from the framework, i.e., $G_{i+1} = G_i \cup (X_i \setminus \mathcal{A}^?)$, $\mathcal{A}_{i+1}^? = \mathcal{A}_i^? \setminus X_i$, and $\mathcal{R}_{i+1} = \mathcal{R}_i|_{\mathcal{A} \cup \mathcal{A}_{i+1}^?}$. Set $i \leftarrow i + 1$.
4. Repeat the previous step until $G_i = G_{i-1}$.
5. The ungrounded completion of IAF for S is $IAF_S^{\text{ungr}} = \langle \mathcal{A}_S^{\text{ungr}}, \mathcal{R}|_{\mathcal{A}_S^{\text{ungr}}} \rangle$ with $\mathcal{A}_S^{\text{ungr}} = \mathcal{A} \cup \mathcal{A}_i^?$.

Intuitively, the ungrounded completion removes all and only those arguments that are in S and that are possible candidates for membership in the grounded extension (elements of X_i in each iteration i)—all other arguments are included. The purpose of that is to make it as unlikely as possible for S to be grounded in this completion. The ungrounded completion is *critical* in the following sense: If a necessarily complete set S is grounded even in the ungrounded completion, then it must be grounded in all completions. This is formalized in Lemma 16.

Lemma 16. *Let $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete argumentation framework, $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ be a necessarily complete set of arguments in IAF, and let IAF_S^{ungr} be the ungrounded completion of IAF for S. S is the necessarily grounded extension of IAF if and only if $S|_{\mathcal{A}_S^{\text{ungr}}}$ is the grounded extension of IAF_S^{ungr}.*

Proof. If $S|_{\mathcal{A}_S^{\text{ungr}}}$ is not the grounded extension of IAF_S^{ungr}, it immediately follows that S is not necessarily grounded in IAF. We now prove the other direction of the equivalence: Let $S|_{\mathcal{A}_S^{\text{ungr}}}$ be the grounded extension of IAF_S^{ungr}. We prove that, then, S is necessarily grounded in IAF.

First, we observe that whenever $S|_{\mathcal{A}_S^{\text{ungr}}}$ is the grounded extension of IAF_S^{ungr} (which we know by assumption), then $S|_{\mathcal{A}_S^{\text{ungr}}} = G_{i'}$ for the set $G_{i'}$ in the last iteration i' of the algorithm: $G_{i'} \subseteq S|_{\mathcal{A}_S^{\text{ungr}}}$ holds because, by construction, $G_{i'}$ consists only of definite arguments, and $S|_{\mathcal{A}_S^{\text{ungr}}} \subseteq G_{i'}$ holds because $S|_{\mathcal{A}_S^{\text{ungr}}}$ is grounded in IAF_S^{ungr} and no argument outside of $G_{i'}$ could be acceptable with respect to $G_{i'}$ in the ungrounded completion. Since $G_{i'}$ consists only of definite arguments, we know that $S|_{\mathcal{A}_S^{\text{ungr}}}$ consists only of definite arguments under the given assumptions.

Now, let $IAF^* = \langle \mathcal{A}^*, \mathcal{R}|_{\mathcal{A}^*} \rangle$ be any completion of $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ (different from the ungrounded completion) and let G^* be its grounded extension. Since we know by assumption that $S|_{\mathcal{A}^*}$ is complete in IAF^*, with the fact (proven by Dung (1995)) that the grounded extension is contained in all complete extensions of the same argumentation framework, we can conclude that $G^* \subseteq S|_{\mathcal{A}^*}$.

However, we also have $S|_{\mathcal{A}^*} \subseteq G^*$: Since $S|_{\mathcal{A}_S^{\text{ungr}}}$ contains only definite arguments, these must be in G^*, too. Now assume that $S|_{\mathcal{A}^*} \not\subseteq G^*$. Then there is a possible (non-definite) argument $a \in (S|_{\mathcal{A}^*} \setminus G^*)$. We know that a is not included in the ungrounded completion. We also know that a is not acceptable with respect to G^* in IAF^*, because otherwise it would need to be included in the grounded set G^*. Also, since $S|_{\mathcal{A}_S^{\text{ungr}}} \subseteq G^*$, a is not acceptable with respect to $S|_{\mathcal{A}_S^{\text{ungr}}}$ either (remember that S is necessarily complete and, in particular, necessarily conflict-free in IAF, so any attackers must be outside of S). So, there must be an attacker $b \notin S$ of a which is not attacked by G^* (and, therefore, not attacked by $S|_{\mathcal{A}_S^{\text{ungr}}}$) in IAF^*. Since the ungrounded completion includes all arguments that are not in S, b is also included in $\mathcal{A}_S^{\text{ungr}}$. Further, since the ungrounded completion includes all and only those possible attacks that target S, the attack (b,a) is included and any possible defending attacks are not included in the ungrounded completion. However, this means that the attack (b,a) is not defended by $S|_{\mathcal{A}_S^{\text{ungr}}}$ in the ungrounded completion, which, by its construction, would mean that a would be included in $\mathcal{A}_S^{\text{ungr}}$ (a could only be excluded in Step 3 if it is acceptable with respect to a subset of $S|_{\mathcal{A}_S^{\text{ungr}}}$, which a is not, due to the attack by b). This contradicts the fact that a is not included in the ungrounded completion. Therefore, such an argument a cannot exist and we can conclude $S|_{\mathcal{A}^*} \subseteq G^*$ and, in total, $S|_{\mathcal{A}^*} = G^*$. So, $S|_{\mathcal{A}^*}$ is grounded in IAF^* and, since IAF^* was kept generic, S is necessarily grounded in IAF. □

Theorem 17. GR-INCNV *is in* P.

Proof. Let $(\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle, S)$ be an instance of GR-INCNV. If S is not necessarily complete in $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, it is not necessarily grounded in $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, either. By Theorem 13, the former can be checked in polynomial time. Therefore, we may assume that S is necessarily complete. Given a completion, GR-VERIFICATION can be solved in polynomial time, and Lemma 16 yields that the answer to GR-INCNV is the same as that to GR-VERIFICATION for the ungrounded completion. Since the ungrounded completion can clearly be constructed in polynomial time, this completes the proof. □

Again, the P upper bound immediately transfers to the argument-incomplete model.

Corollary 18. GR-ARGINCNV *is in* P.

We have completed our proofs for P membership of necessary verification in all three incompleteness models for the admissible, stable, complete, and grounded semantics.

4.2 Lower Bounds

Our final results show that the complexity of possible verification for the preferred semantics raises from coNP-hardness to Σ_2^p-completeness in all three models.

Theorem 19. PR-ATTINCPV *is* Σ_2^p*-hard.*

Proof. First, we quickly recall some notation from propositional logic. A boolean variable x has two literals, x and $\neg x$. A boolean formula is in conjunctive normal form (CNF) if it is a conjunction of disjunctions of literals (clauses), and in disjunctive normal form (DNF) if it is a disjunction of conjunctive clauses of literals. 3-CNF (respectively, 3-DNF) denotes CNF (respectively, DNF) with at most three literals per clause. A truth assignment τ on a set X of variables is a function $\tau : X \to \{\mathit{true}, \mathit{false}\}$. For a formula φ and truth assignments $\tau_1, \tau_2, \ldots, \tau_k$ on disjoint sets of variables, $\varphi[\tau_1, \tau_2, \ldots, \tau_k]$ denotes the formula obtained by replacing variables in φ with their truth values in $\tau_1, \tau_2, \ldots, \tau_k$.

To prove Σ_2^p-hardness, we reduce from the quantified satisfiability problem Σ_2SAT, which is well-known to be complete for Σ_2^p (Stockmeyer 1976): Given a 3-DNF formula φ on two disjoint sets of variables, X and Y, does $\exists \tau_X \, \forall \tau_Y : \varphi[\tau_X, \tau_Y]$ evaluate to true (where τ_X and τ_Y are truth assignments on X and Y, respectively)?

Let (φ, X, Y) be an instance of Σ_2SAT, where $X = \{x_1, \ldots, x_{|X|}\}$ and $Y = \{y_1, \ldots, y_{|Y|}\}$ are two disjoint sets of propositional variables and φ is a 3-DNF formula over $X \cup Y$. For $\bar{\varphi} = \neg\varphi$, the question in Σ_2SAT is equivalent to asking whether $\exists \tau_X \forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \mathit{false}$, where $\bar{\varphi} = c_1 \wedge \cdots \wedge c_m$ is a formula in 3-CNF with clauses c_1 through c_m. From now on, we will mostly use this CNF formulation of the problem.

We create an instance $(\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle, S)$ of PR-ATTINCPV from (φ, X, Y) as follows (see Figure 3a for an example):

$$\mathcal{A} = \left\{ \begin{array}{ll} y_i, \bar{y_i}, & \text{for } y_i \in Y \\ x_i, \bar{x_i}, & \text{for } x_i \in X \\ c_i, & \text{for } c_i \text{ in } \bar{\varphi} \\ s & \end{array} \right\},$$

$$\mathcal{R}^? = \{ \ (s, \bar{x_i}), \quad \text{for } x_i \in X \ \},$$

$$\mathcal{R} = \left\{ \begin{array}{ll} (\bar{y_i}, y_i), (y_i, \bar{y_i}), & \text{for } y_i \in Y \\ (\bar{x_i}, x_i), & \text{for } x_i \in X \\ (c_i, c_i), & \text{for } c_i \text{ in } \bar{\varphi} \\ (c_i, y_j), (c_i, \bar{y_j}), & \text{for } c_i \text{ in } \bar{\varphi}, y_j \in Y \\ (c_i, x_k), (c_i, \bar{x_k}), & \text{for } c_i \text{ in } \bar{\varphi}, x_k \in X \\ (y_j, c_i), & \text{if } y_j \text{ in } c_i \\ (\bar{y_j}, c_i), & \text{if } \neg y_j \text{ in } c_i \\ (x_k, c_i), & \text{if } x_k \text{ in } c_i \\ (\bar{x_k}, c_i), & \text{if } \neg x_k \text{ in } c_i \end{array} \right\}.$$

Finally, let $S = \{s\}$. We call all arguments x_i, $\bar{x_i}$, y_i, and $\bar{y_i}$ *literal arguments* and arguments c_i *clause arguments*. Note that S is necessarily admissible in $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$, so the verification of possible preferredness boils down to checking whether all supersets of S are nonadmissible in some completion of $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$.

We prove that $(\varphi, X, Y) \in \Sigma_2\text{SAT} \Leftrightarrow (\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle, S) \in$ PR-ATTINCPV. Assume that $(\varphi, X, Y) \in \Sigma_2$SAT, i.e., $\exists \tau_X \forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \mathit{false}$. Let τ_X be an assignment of truth values to the variables in X that satisfies $\forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \mathit{false}$. Let $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$ be the completion of $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$ obtained by letting $\mathcal{R}^{\tau_X} = \mathcal{R} \cup \{(s, \bar{x_i}) \in \mathcal{R}^? \mid \tau_X(x_i) = \mathit{true}\}$. In $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$, the assignment τ_X to the variables in X is translated to a commitment on literal arguments: If, for $x_i \in X$, $\tau_X(x_i) = \mathit{true}$, then the attack by s against argument $\bar{x_i}$ is included and $\bar{x_i}$ can no longer be a member of admissible supersets of S, while argument x_i is defended by s and potentially can be such a member. On the other hand, if $\tau_X(x_i) = \mathit{false}$, the attack is excluded and the roles are switched: Argument x_i cannot be defended against argument $\bar{x_i}$ by S (or any conflict-free superset of S), so x_i cannot be contained in admissible supersets of S, whereas $\bar{x_i}$ can.

Now let τ_Y be any truth assignment for Y. We know that $\bar{\varphi}[\tau_X, \tau_Y] = \mathit{false}$. Transform τ_X and τ_Y to a set $S_{(\tau_X, \tau_Y)} \supset S$ of arguments by letting $S_{(\tau_X, \tau_Y)} = S \cup \{x_i \mid \tau_X(x_i) = \mathit{true}\} \cup \{\bar{x_i} \mid \tau_X(x_i) = \mathit{false}\} \cup \{y_i \mid \tau_Y(y_i) = \mathit{true}\} \cup \{\bar{y_i} \mid \tau_Y(y_i) = \mathit{false}\}$. It is easy to see that $S_{(\tau_X, \tau_Y)}$ is conflict-free in $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$. However, $S_{(\tau_X, \tau_Y)}$ cannot defend itself against all clause arguments $c_1, \ldots, c_m$ in $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$, and therefore is not admissible: Since $\bar{\varphi}$ is in CNF and $\bar{\varphi}[\tau_X, \tau_Y] = \mathit{false}$, at least one clause in $\bar{\varphi}$ is unfulfilled. Let c_j be any such clause. Since the clauses of $\bar{\varphi}$

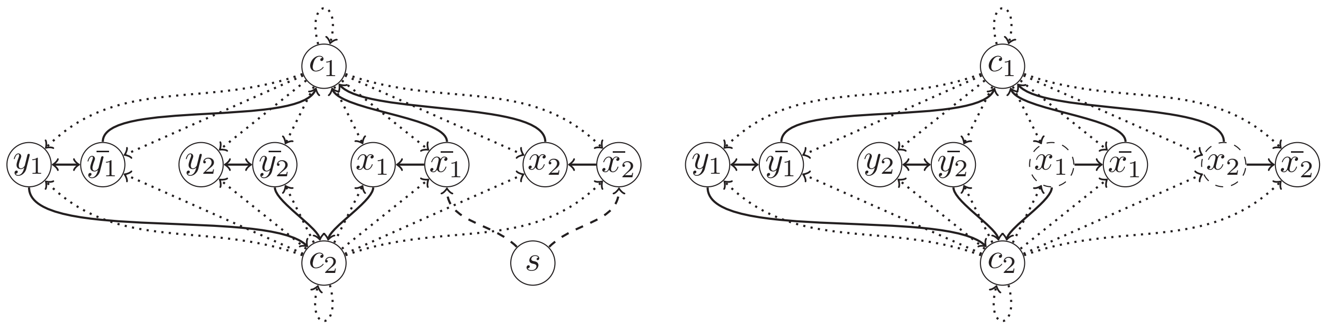

(a) Graph representation of $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$

(b) Graph representation of $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$

Figure 3: Graph representations created from clauses $c_1 = (\neg x_1 \vee x_2 \vee \neg y_1)$ and $c_2 = (x_1 \vee y_1 \vee \neg y_2)$, where dashed attacks or arguments indicate uncertainty as usual, and attacks by clause arguments are displayed as dotted arcs to facilitate readability.

are disjunctions of literals, all literals in c_j are unfulfilled. The only arguments in $\mathcal{A}$ that attack the clause argument c_j are the literal arguments whose corresponding literals appear in clause c_j. However, by construction, none of these arguments are in $S_{(\tau_X,\tau_Y)}$, since all these literals are *false* in τ_X and τ_Y. Therefore, no argument in $S_{(\tau_X,\tau_Y)}$ attacks argument c_j. On the other hand, c_j attacks all literal arguments and therefore it attacks $S_{(\tau_X,\tau_Y)}$, which proves that $S_{(\tau_X,\tau_Y)}$ is not admissible in $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$. All other supersets of S are either a subset of $S_{(\tau_X,\tau_Y)}$ or not conflict-free, and thus can't be admissible, either. Since τ_Y was kept generic, this covers all possible supersets of S and proves that S is preferred in $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$, and we have $(\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle, S) \in$ PR-ATTINCPV.

For the other direction, assume that $(\varphi, X, Y) \notin \Sigma_2\mathrm{SAT}$, i.e., $\forall \tau_X \exists \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = true$. Let τ_X be any assignment on X and let τ_Y be an assignment on Y that satisfies $\bar{\varphi}[\tau_X, \tau_Y] = true$. Create the completion $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$ and the set $S_{(\tau_X,\tau_Y)}$ as before. Since $\bar{\varphi}[\tau_X, \tau_Y] = true$, all clauses in $\bar{\varphi}$ are fulfilled, which means that in each clause at least one literal must be fulfilled. Each such literal corresponds to a literal argument in $S_{(\tau_X,\tau_Y)}$, which attacks the corresponding clause argument. So, $S_{(\tau_X,\tau_Y)}$ is admissible, which shows that S is not preferred in $\langle \mathcal{A}, \mathcal{R}^{\tau_X} \rangle$, and since τ_X was generic, S is not preferred in any completion of $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$, which proves $(\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle, S) \notin$ PR-ATTINCPV. □

The same hardness can be proven for the argument-incomplete model. We omit the proof due to space limitations. Figure 3b and Example 21 may give a rough idea, though.

Theorem 20. PR-ARGINCPV *is* Σ_2^p*-hard.*

Example 21. Consider a $\Sigma_2\mathrm{SAT}$ instance (φ, X, Y) with $X = \{x_1, x_2\}$, $Y = \{y_1, y_2\}$ and $\varphi = (x_1 \wedge \neg x_2 \wedge y_1) \vee (\neg x_1 \wedge \neg y_1 \wedge y_2)$. We have $\bar{\varphi} = \neg\varphi = c_1 \wedge c_2$ with $c_1 = (\neg x_1 \vee x_2 \vee \neg y_1)$ and $c_2 = (x_1 \vee y_1 \vee \neg y_2)$. We have $(\varphi, X, Y) \notin \Sigma_2\mathrm{SAT}$, because for all assignments τ_X on X and the assignment τ_Y with $\tau_Y(y_1) = false$, $\tau_Y(y_2) = false$ we have $\varphi[\tau_X, \tau_Y] = false$; equivalently, $\bar{\varphi}[\tau_X, \tau_Y] = true$.

Figures 3a and 3b show, respectively, the graph representations of the incomplete argumentation frameworks in the instances $(\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle, \{s\})$ and $(\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle, \emptyset)$ that are created from (φ, X, Y) according to the constructions in the proofs of Theorems 19 and 20. Attacks by clause arguments are displayed as dotted arcs to facilitate readability. Both instances are NO-instances for PR-ATTINCPV and PR-ARGINCPV, respectively. The set $\{s, \bar{y_1}, \bar{y_2}\}$ (corresponding to τ_Y from above) is an admissible superset of $\{s\}$ in all completions of $\langle \mathcal{A}, \mathcal{R}, \mathcal{R}^? \rangle$, while the set $\{\bar{y_1}, \bar{y_2}\}$ is an admissible superset of $\emptyset$ in all completions of $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R} \rangle$.

Both previous results also provide Σ_2^p-hardness for the problem PR-INCPV in the general model, which completes our complexity analysis.

Corollary 22. PR-INCPV *is* Σ_2^p*-hard.*

5 Conclusion

The complexity results show a pattern in how introducing incomplete information affects the complexity of the verification problem in abstract argumentation frameworks. We observe that there are only two triggers for an increase of complexity: the preferred semantics for possible verification in all three models, and the admissible semantics (along with all other semantics that entail admissibility) for possible verification in the argument-incomplete model (and, therefore, also in the general incomplete model). In all other cases—in particular, for all variants of necessary verification—introducing incomplete information does not make the verification problem computationally harder. Table 1 gives an overview. Note that each of our hardness results for verification problems carries over to any generalized model; so our approach is potentially useful also in other frameworks.

A task for future work is to analyze the parameterized complexity of the problems studied here, the complexity of possible and necessary variants of other decision problems than verification, and to look at further semantics.

Acknowledgements

We thank the reviewers for their helpful comments. This work was supported in part by a grant from the NRW Ministry for Innovation, Science, and Research.

Table 1: Summary of the complexity results, where "$\mathcal{C}$-c." denotes completeness for a complexity class $\mathcal{C}$. All previously known results are attributed to the respective source, and all other results are novel.

s	VER	ATTINCPV	ATTINCNV	ARGINCPV	ARGINCNV	INCPV	INCNV
CF	in P♠	in P★	in P★	in P♦	in P♦	in P	in P
AD	in P♠	in P▲	in P★	NP-c.♦	in P	NP-c.	in P
ST	in P♠	in P▲	in P▲	NP-c.♦	in P	NP-c.	in P
CP	in P♠	in P▲	in P▲	NP-c.♦	in P	NP-c.	in P
GR	in P♠	in P▲	in P▲	NP-c.♦	in P	NP-c.	in P
PR	coNP-c.♣	Σ_2^p-c.	coNP-c.▲	Σ_2^p-c.	coNP-c.♦	Σ_2^p-c.	coNP-c.

♠ (Dung 1995), ♣ (Dimopoulos and Torres 1996), ★ (Coste-Marquis et al. 2007),
▲ (Baumeister, Neugebauer, and Rothe 2015), ♦ (Baumeister, Rothe, and Schadrack 2015)

References

Baumann, R., and Brewka, G. 2010. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proceedings of the 3rd International Conference on Computational Models of Argument*, 75–86. IOS Press.

Baumeister, D.; Neugebauer, D.; and Rothe, J. 2015. Verification in attack-incomplete argumentation frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, 341–358. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*.

Baumeister, D.; Rothe, J.; and Schadrack, H. 2015. Verification in argument-incomplete argumentation frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, 359–376. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*.

Boella, G.; Kaci, S.; and van der Torre, L. 2009. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Proceedings of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 107–118. Springer-Verlag *Lecture Notes in Artificial Intelligence #5590*.

Cayrol, C.; de Saint-Cyr, F.; and Lagasquie-Schiex, M. 2010. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research* 38:49–84.

Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.; and Marquis, P. 2007. On the merging of Dung's argumentation systems. *Artificial Intelligence* 171(10):730–753.

Coste-Marquis, S.; Konieczny, S.; Mailly, J.; and Marquis, P. 2015. Extension enforcement in abstract argumentation as an optimization problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2876–2882. AAAI Press/IJCAI.

Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science* 170(1):209–244.

Dung, P., and Thang, P. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proceedings of the 3rd International Conference on Computational Models of Argument*, 171–182. IOS Press.

Dung, P. 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., and Wooldridge, M. 2009. Complexity of abstract argumentation. In Rahwan, I., and Simari, G., eds., *Argumentation in Artificial Intelligence*. Springer. chapter 5, 85–104.

Hunter, A. 2014. Probabilistic qualification of attack in abstract argumentation. *International Journal of Approximate Reasoning* 55(2):607–638.

Janssen, J.; Cock, M. D.; and Vermeir, D. 2008. Fuzzy argumentation frameworks. In *Proceedings of the 12th International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems*, 513–520.

Li, H.; Oren, N.; and Norman, T. 2011. Probabilistic argumentation frameworks. In *Proceedings of the 1st International Workshop on Theory and Applications of Formal Argumentation*, 1–16. Springer-Verlag *Lecture Notes in Artificial Intelligence #7132*.

Liao, B.; Jin, L.; and Koons, R. 2011. Dynamics of argumentation systems: A division-based method. *Artificial Intelligence* 175(11):1790–1814.

Maher, M. 2016. Resistance to corruption of strategic argumentation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 1030–1036. AAAI Press.

Papadimitriou, C. 1995. *Computational Complexity*. Addison-Wesley, 2nd edition. Reprinted with corrections.

Rienstra, T. 2012. Towards a probabilistic Dung-style argumentation system. In *Proceedings of the 1st International Conference on Agreement Technologies*, 138–152. CEUR Workshop Proceedings.

Rothe, J. 2005. *Complexity Theory and Cryptology. An Introduction to Cryptocomplexity*. EATCS Texts in Theoretical Computer Science. Springer-Verlag.

Stockmeyer, L. 1976. The polynomial-time hierarchy. *Theoretical Computer Science* 3(1):1–22.

Wallner, J.; Niskanen, A.; and Järvisalo, M. 2016. Complexity results and algorithms for extension enforcement in abstract argumentation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 1088–1094. AAAI Press.