Complexity of Verification in Incomplete Argumentation Frameworks

Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack

Abstract

Abstract argumentation frameworks are a well-established formalism to model nonmonotonic reasoning processes. However, Dung's model [19] cannot express incomplete or conflicting knowledge about the state of a given argumentation. In previous work [8, 6, 7] we considered incomplete argumentation frameworks which allow uncertainty regarding the set of attacks, the set of arguments, or both. Different semantics are used in order to identify sets of strong arguments. An important task is to verify whether a given set of arguments fulfills the criteria for a given semantics. We complement existing results on the complexity of variants of the verification problem in incomplete argumentation frameworks and provide a full complexity map covering all three models and all classical semantics.

1 Introduction

Within the field of artificial intelligence, abstract argumentation frameworks have emerged as a prominent 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 [19] 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 about the existence of elements in an argumentation framework. Such *qualitative* uncertainty about the state of an argumentation framework was introduced by Coste-Marquis et al. [16] for the set of attacks. Baumeister et al. [8] propose an extended model that allows uncertainty about the set of arguments, or about both attacks and arguments. In applications, such incomplete argumentation frameworks may arise as intermediate states in an elicitation process, or when merging or aggregating different beliefs (i.e., the agents' individual, subjective views) about an argumentation framework's state, or in cases where complete information cannot be obtained. The main goal of this paper 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). This approach has already been taken in various areas of computational social choice: in voting by, e.g., Konczak and Lang [25], Xia and Conitzer [36], Chevaleyre et al. [15], and Baumeister et al. [3, 4]; in fair division by Bouveret et al. [12] and Baumeister et al. [9]; in algorithmic game theory by Lang et al. [26]; and in judgment aggregation by Baumeister et al. [5]. However, this approach is new to argumentation theory: In two of this paper's predecessors, Baumeister et al. [6, 7] were the first to define and study possible and necessary verification for certain semantics in incomplete argumentation frameworks, and they continued this line of research in their recent work [10]. The present paper merges and extends these preliminary versions.

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 [24] 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 et al. [27] assume that a probability distribution over both arguments and attacks is given. Other approaches associate a probability with each set of arguments [20, 31] to indicate whether all and only these arguments are active, or with each spanning subtree of the argument graph [23] 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 [14, 11], the complexity of computing the acceptability of a single argument after changing the arguments or attacks [28], or enforcement of a set of arguments [2, 35, 17], where the question is how much a given argumentation framework needs to be modified to make the given set of arguments acceptable. Maher [29] studies a strategic version of enforcement, focusing on resistance to corruption.

In the following, 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 future tasks (Section 5).

2 Preliminaries

We start by defining argumentation frameworks due to Dung [19], mostly following the notation by Dunne and Wooldridge [22].

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

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

Example 2. Figure 1 displays the graph representation of the argumentation framework $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ with $\mathscr{A} = \{a,b,c\}$ and $\mathscr{R} = \{(c,a),(c,b)\}$.

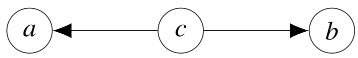

Figure 1: A simple argumentation framework

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 [19]. They are formalized in Definition 3, after introducing some necessary notions.

An argument $a \in \mathscr{A}$ is *defended by* $S \subseteq \mathscr{A}$ if, for each $b \in \mathscr{A}$ with $(b,a) \in \mathscr{R}$, there is a $c \in S$ such that $(c,b) \in \mathscr{R}$. For an argumentation framework AF, the *characteristic function* $F_{AF} : 2^{\mathscr{A}} \to 2^{\mathscr{A}}$ maps each set S of arguments to the set of arguments that are defended by S, i.e., $F_{AF}(S) = \{a \in \mathscr{A} \mid a \text{ is defended by } S\}$. The characteristic function always has a least fixed point, since it is monotonic with respect to set inclusion. Let F_{AF}^k 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 $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ be an argumentation framework. A set $S \subseteq \mathscr{A}$ is *conflict-free* if $(a,b) \notin \mathscr{R}$ for all $a,b \in S$. A conflict-free set $S \subseteq \mathscr{A}$ 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 $b \in \mathscr{A} \setminus S$ there is an $a \in S$ with $(a,b) \in \mathscr{R}$.

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 [19] 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.

Figure 2 displays all relations among the various semantics that we use. If an area labeled with semantics $\mathbf{s}$ is fully included in an area labeled with semantics $\mathbf{s}'$, this indicates that in all argumentation frameworks all sets of arguments that fulfill $\mathbf{s}$ also fulfill $\mathbf{s}'$. The converse is not necessarily true, i.e., all displayed set inclusions are strict. Further, none of the areas are disjoint, so one and the same set of arguments might fulfill all semantics simultaneously.

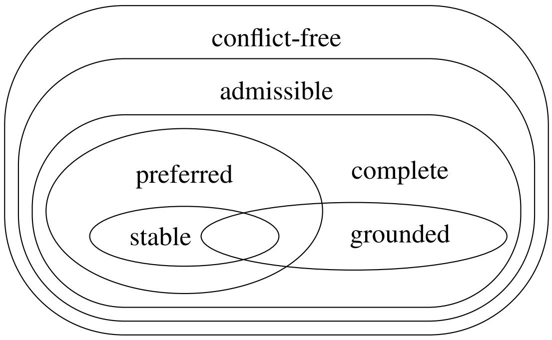

Figure 2: Relations among various semantics for sets of arguments

We assume the reader to be familiar with the complexity classes of the polynomial hierarchy, in particular, P, NP, coNP, and $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$, as well as the concepts of hardness and completeness. For an introduction, see, e.g., the books by Papadimitriou [30] and Rothe [32].

Dunne and Wooldridge [22] defined decision problems regarding the existence or status of acceptable arguments. We focus on the verification problem $\mathbf{s}$-VERIFICATION, which is parameterized by one of the semantics (denoted $\mathbf{s}$) defined above and asks whether for an argumentation framework $\langle \mathscr{A}, \mathscr{R} \rangle$ a given subset of the 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. The problem PR-VERIFICATION was shown to be coNP-complete by Dimopoulos and Torres [18], but Dung [19] 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 $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$, where $\mathscr{A}$ and $\mathscr{A}^?$ are disjoint sets of arguments and $\mathscr{R}$ and $\mathscr{R}^?$ are disjoint subsets of $(\mathscr{A} \cup \mathscr{A}^?) \times (\mathscr{A} \cup \mathscr{A}^?)$. $\mathscr{A}$ is the set of arguments that are known to definitely exist, while $\mathscr{A}^?$ contains all possible additional arguments not (yet) known to exist. Similarly, $\mathscr{R}$ is the set of attacks that are known to definitely exist (as long as both incident arguments turn out to exist), while $\mathscr{R}^?$ contains all possible additional attacks not (yet) known to exist.

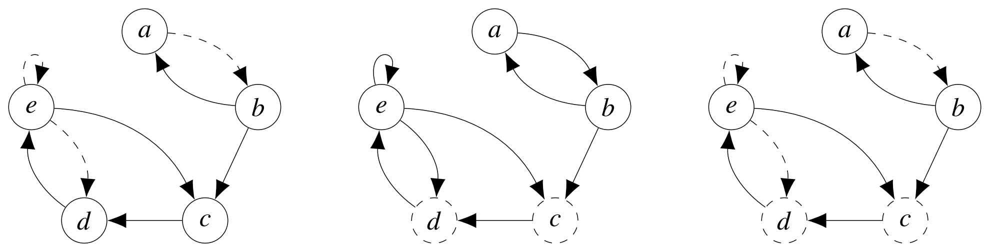

(a) An attack-incomplete AF (b) An argument-incomplete AF (c) An incomplete AF

Figure 3: Some examples of incomplete argumentation frameworks

Example 5. Figure 3 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 3a has no uncertainty regarding the arguments, while the one in Figure 3b has no uncertainty regarding the attacks. The incomplete argumentation framework in Figure 3c combines the uncertainty of the other two.

An incomplete argumentation framework $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ 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 $\mathscr{A}^?$ and each possible attack in $\mathscr{R}^?$ is either included or excluded. Such an argumentation framework is called a *completion of* $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$. When excluding a possible argument, all its incident attacks are also automatically excluded: For a set $\mathscr{A}^*$ of arguments with $\mathscr{A} \subseteq \mathscr{A}^* \subseteq \mathscr{A} \cup \mathscr{A}^?$, the *restriction of a relation* $\mathscr{R}$ *to* $\mathscr{A}^*$ is $\mathscr{R}|_{\mathscr{A}^*} = \{(a,b) \in \mathscr{R} \mid a,b \in \mathscr{A}^*\}$.

Definition 6. Let $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ be an incomplete argumentation framework. An argumentation framework $IAF^* = \langle \mathscr{A}^*, \mathscr{R}^* \rangle$ with $\mathscr{A} \subseteq \mathscr{A}^* \subseteq \mathscr{A} \cup \mathscr{A}^?$ and $\mathscr{R}|_{\mathscr{A}^*} \subseteq \mathscr{R}^* \subseteq \left(\mathscr{R} \cup \mathscr{R}^?\right)|_{\mathscr{A}^*}$ 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 $2^{|\mathscr{R}^?|+|\mathscr{A}^?|}$, 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 3a and 3b have $2^3 = 8$ and $2^2 = 4$ completions, respectively. The incomplete argumentation framework in Figure 3c 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 **s**:

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

In the **s**-INC-NECESSARY-VERIFICATION (**s**-INCNV) problem the input is the same, but the question is whether for all completions $IAF^* = \langle \mathscr{A}^*, \mathscr{R}^* \rangle$ of IAF, $S|_{\mathscr{A}^*} = S \cap \mathscr{A}^*$ is an **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. Note that these definitions of possible and necessary verification allow a subset of S—namely, $S|_{\mathscr{A}^*}$—to be an extension. This follows the intuition that for S to be an extension, no element of S may be unaccepted, but it is not harmful if elements of S are discarded. An alternative definition that strictly requires S to be an extension could be obtained as a special case by restricting to instances where $S \cap \mathscr{A}^? = \emptyset$.

Incomplete argumentation frameworks are a generalization of both pure models of incomplete argumentation frameworks. Fixing $\mathscr{A}^? = \emptyset$ in Definitions 4 and 6 yields exactly the class of attack-incomplete argumentation frameworks as proposed by Coste-Marquis et al. [16] (and further studied by Baumeister et al. [6]), and fixing $\mathscr{R}^? = \emptyset$ yields exactly the class of argument-incomplete argumentation frameworks as proposed by Baumeister et al. [7]. In attack-incomplete argumentation frameworks, the set of possible arguments can be omitted and it can be written as $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$. Likewise, $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ 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 et al [6, 7]. In the attack-incomplete model, we write **s**-ATTINCPV and **s**-ATTINCNV for possible and necessary verification, respectively, and **s**-ARGINCPV and **s**-ARGINCNV in the argument-incomplete model.

4 Complexity of Possible and Necessary Verification

In this section, we complete the complexity analysis of possible and necessary verification in all three presented models of incompleteness and for the conflict-free, admissible, stable, complete, grounded, and preferred semantics. All results are summarized in Table 1 in Section 5. Since general incomplete argumentation frameworks are a generalization of both individual models of incompleteness, all upper complexity bounds for the general model carry over to both individual models, and all lower complexity bounds for any of the individual models carry over to the general model.

4.1 Upper Bounds

We start by a simple Σ_2^p upper bound for PR-INCPV.

Theorem 8. PR-INCPV *is in* Σ_2^p.

Proof. The result follows directly from the quantifier representation of the problem. 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. In the case of PR-INCPV, this polynomial-time predicate is preceded first by an existential quantifier (guessing a completion) and then a universal quantifier (verifying preferredness) yielding Σ_2^p membership. □

We turn to proving P membership for the remaining open problems, starting with AD-INCNV and ST-INCNV.

Theorem 9. AD-INCNV *and* ST-INCNV *both are in* P.

Proof. Let (IAF,S) with $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ be an instance of AD-INCNV. Let $IAF_S^{pes} = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}_S^{pes} \rangle$ with $\mathscr{R}_S^{pes} = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$ be the pessimistic argument-incomplete argumentation framework obtained when eliminating attack incompleteness by including each and only those attacks that target S (which can clearly be done in polynomial time).

We will prove that $(IAF,S) \in$ AD-INCNV $\iff (IAF_S^{pes},S) \in$ AD-ARGINCNV. Since AD-ARGINCNV $\in$ P and IAF_S^{pes} can be created from IAF in polynomial time, this yields that AD-INCNV $\in$ P. A completely analogous argument applies to the stable semantics and the problem ST-INCNV.

If $(IAF,S) \in$ AD-INCNV, then $(IAF_S^{pes},S) \in$ AD-ARGINCNV follows trivially, since the set of completions of IAF_S^{pes} is a subset of the completions of IAF. We prove the other direction of the equivalence by contraposition. Assume that $(IAF,S) \notin$ AD-INCNV. Then there is a completion IAF^* of IAF in which S is not admissible. Create a completion IAF_S^{pes*} from the argument-incomplete argumentation framework IAF_S^{pes} by adding exactly those elements of $\mathscr{A}^?$ to the set of arguments that are also added in IAF^*. By construction, in IAF_S^{pes*} all attacks against arguments in S that exist in IAF^* are included, too, and any attacks against arguments outside of S that are not in IAF^* are not included, either. Since S is not admissible in IAF^*, it can clearly not be admissible in IAF_S^{pes*}. Therefore, we have $(IAF_S^{pes},S) \notin$ AD-ARGINCNV. This completes the proof. ❑

Turning to the complete and grounded semantics, we can successively prove P membership of CP-INCNV and GR-INCNV in Theorems 10 and 15, respectively.

Theorem 10. CP-INCNV *is in* P.

Proof. Let (IAF,S) with $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ be an instance of CP-INCNV. Since AD-INCNV $\in$ P, we may assume that S is necessarily admissible in IAF. Then, 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 criterion.

If all arguments $a \in (\mathscr{A} \cup \mathscr{A}^?) \setminus S$ are definitely attacked by S, i.e., $(b_a,a) \in \mathscr{R}$ 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 $a \in (\mathscr{A} \cup \mathscr{A}^?) \setminus S$ be any argument outside of S that is not definitely attacked by S, i.e., $(b,a) \notin \mathscr{R}$ for all $b \in S \cap \mathscr{A}$ (if a were attacked by S, it clearly could not be defended by S in any completion). Let $Att(a) = \{c \in \mathscr{A} \cup \mathscr{A}^? \mid (c,a) \in \mathscr{R}\}$ be the set of all arguments with a definite attack against a. Further, let $\mathscr{R}_a = \mathscr{R} \cup \{(s,c) \in \mathscr{R}^? \mid s \in S \text{ and } c \in Att(a) \setminus \{a\}\}$ 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 $C_a = \langle \mathscr{A}_a, \mathscr{R}_a|_{\mathscr{A}_a} \rangle$ where $\mathscr{A}_a = \mathscr{A} \cup \{a\} \cup \{d \in \mathscr{A}^? \mid (d,a) \notin \mathscr{R}_a\}$, i.e., C_a uses the attack relation $\mathscr{R}_a$ and includes a and exactly those possible arguments that do not attack a (in $\mathscr{R}_a$). 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 $d \in \mathscr{A}_a$ with $(d,a) \in \mathscr{R}_a|_{\mathscr{A}_a}$ and S does not attack d in C_a. By construction of C_a, we know that d is a definite argument, i.e., $d \in \mathscr{A}$, and (d,a) is a definite attack, i.e., $(d,a) \in \mathscr{R}$, so d attacks a in any completion that contains a. Also, in all completions S either does not defend a against d, 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. ❑

The following upper bound then follows immediately.

Corollary 11. 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 12. Let $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ be an incomplete argumentation framework and $S \subseteq \mathscr{A} \cup \mathscr{A}^?$ be a set of arguments in IAF. The *ungrounded completion* IAF_S^{ungr} of IAF for S is the completion that is obtained by the following algorithm. The algorithm first eliminates attack incompleteness and then defines a finite sequence $(IAF_i)_{i \geq 0}$ of argument incomplete argumentation frameworks, with the ungrounded completion being the maximal completion (that includes all remaining possible arguments) of the sequence's last element.

1. Eliminate attack incompleteness: Let $\mathscr{R}_0 = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$, i.e., include only those possible attacks that attack S.
2. Let initially $G_0 = \emptyset$, $\mathscr{A}_0^? = \mathscr{A}^?$, $IAF_0 = \langle \mathscr{A}, \mathscr{A}_0^?, \mathscr{R}_0 \rangle$ 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., $X_i = F_{Max_i}(G_i) \cap S$. Add the definite arguments in X_i to G and exclude the possible arguments in X_i from the framework, i.e., $G_{i+1} = G_i \cup (X_i \setminus \mathscr{A}^?)$, $\mathscr{A}_{i+1}^? = \mathscr{A}_i^? \setminus X_i$, and $\mathscr{R}_{i+1} = \mathscr{R}_i|_{\mathscr{A} \cup \mathscr{A}_{i+1}^?}$. 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 $IAF_S^{ungr} = \langle \mathscr{A}_S^{ungr}, \mathscr{R}|_{\mathscr{A}_S^{ungr}} \rangle$ with $\mathscr{A}_S^{ungr} = \mathscr{A} \cup \mathscr{A}_i^?$.

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.

Lemma 13 establishes that the ungrounded completion is polynomial-time computable.

Lemma 13. *For an incomplete argumentation framework* $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ *and a set* $S \subseteq \mathscr{A} \cup \mathscr{A}^?$ *of arguments, the ungrounded completion* IAF_S^{ungr} *can be constructed in polynomial time.*

Proof. All individual steps can obviously be carried out in time polynomial in the number of arguments. Also, the loop in Step 4 runs at most a polynomial number of times, since in each execution of the loop there is either (at least) one definite argument that is added to G_{i+1}, or no action is taken in which case the loop terminates. Therefore, the number of times the loop is executed is bounded by the number of definite arguments in the incomplete argumentation framework $AtIAF$. This completes the proof. ❑

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 14. The proof of Lemma 14 is deferred to the appendix.

Lemma 14. *Let* $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ *be an incomplete argumentation framework,* $S \subseteq \mathscr{A} \cup \mathscr{A}^?$ *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* $S|_{\mathscr{A}_S^{ungr}}$ *is the grounded extension of* IAF_S^{ungr}.

Using the above lemmas, we are now ready to show that for the grounded semantics, necessary verification in incomplete argumentation frameworks remains efficient.

Theorem 15. GR-INCNV *is in* P.

Proof. Let $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle, S)$ be an instance of GR-INCNV. If S is not necessarily complete in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$, it is not necessarily grounded in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$, either. By Theorem 10, the former can be checked in polynomial time. Therefore, we may assume that S is necessarily complete.

Lemma 13 provides polynomial-time constructability for the ungrounded completion. Given a completion, GR-VERIFICATION can be solved in polynomial time, and Lemma 14 yields that the answer to GR-INCNV is the same as that to GR-VERIFICATION for the ungrounded completion. ❑

The following upper bound then follows immediately.

Corollary 16. 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

In this section, we prove tight lower bounds for all remaining cases. 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. The reductions used in the proofs of Theorems 17 and 18 are illustrated in Example 19.

Theorem 17. 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 $\tau : X \to \{\texttt{true}, \texttt{false}\}$. 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 (see [33]): Given a 3-DNF formula φ on two disjoint sets of variables, X and Y, the question is whether $\exists \tau_X \, \forall \tau_Y : \varphi[\tau_X, \tau_Y]$ evaluates to $\texttt{true}$ (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 $\exists \tau_X \forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$, 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 $(\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle, S)$ of PR-ATTINCPV from (φ, X, Y) as follows (see Figure 4 for an example):

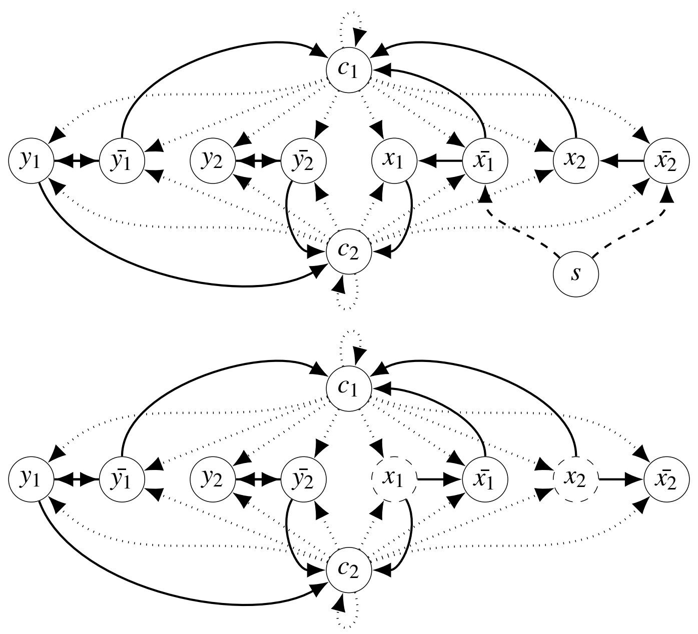

Figure 4: No-instances: Graph representations of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ (top) and $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ (bottom) created from clauses $c_1 = (\neg x_1 \vee x_2 \vee \neg y_1)$ and $c_2 = (x_1 \vee y_1 \vee \neg y_2)$, Dashed attacks or arguments indicate uncertainty and attacks by clause arguments are displayed as dotted arcs to facilitate readability.

$$\mathscr{A} = \left\{ \begin{array}{ll} y_i, \bar{y}_i, & \text{for } y_i \in Y \\ x_i, \bar{x}_i, & \text{for } x_i \in X \\ c_i, & \text{for } c_i \text{ in } \bar{\varphi} \\ s & \end{array} \right\}, \qquad \mathscr{R} = \left\{ \begin{array}{ll} (\bar{y}_i, y_i), (y_i, \bar{y}_i), & \text{for } y_i \in Y \\ (\bar{x}_i, x_i), & \text{for } x_i \in X \\ (c_i, c_i), & \text{for } c_i \text{ in } \bar{\varphi} \\ (c_i, y_j), (c_i, \bar{y}_j), & \text{for } c_i \text{ in } \bar{\varphi}, y_j \in Y \\ (c_i, x_k), (c_i, \bar{x}_k), & \text{for } c_i \text{ in } \bar{\varphi}, x_k \in X \\ (y_j, c_i), & \text{if } y_j \text{ in } c_i \\ (\bar{y}_j, c_i), & \text{if } \neg y_j \text{ in } c_i \\ (x_k, c_i), & \text{if } x_k \text{ in } c_i \\ (\bar{x}_k, c_i), & \text{if } \neg x_k \text{ in } c_i \end{array} \right\}.$$

$$\mathscr{R}^? = \left\{ \; (s, \bar{x}_i), \quad \text{for } x_i \in X \; \right\}.$$

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 $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$, so the verification of possible preferredness boils down to checking whether all supersets of S are nonadmissible in some completion of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$.

We prove that: $(\varphi, X, Y) \in \Sigma_2\text{SAT} \iff (\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle, S) \in \text{PR-ATTINCPV}$.

Assume that $(\varphi, X, Y) \in \Sigma_2\text{SAT}$, i.e., $\exists \tau_X \forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$. Let τ_X be an assignment of truth values to the variables in X that satisfies $\forall \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$. Let $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$ be the completion of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ obtained by letting $\mathscr{R}^{\tau_X} = \mathscr{R} \cup \{(s, \bar{x}_i) \in \mathscr{R}^? \mid \tau_X(x_i) = \texttt{true}\}$. In $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$, the assignment τ_X to the variables in X is translated to a commitment on literal arguments: If, for $x_i \in X$, $\tau_X(x_i) = \texttt{true}$, 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 $\tau_X(x_i) = \texttt{false}$, 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 $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$. Transform τ_X and τ_Y to a set $S_{(\tau_X,\tau_Y)} \supset S$ of arguments by letting

$$\begin{aligned} S_{(\tau_X,\tau_Y)} &= S \cup \{x_i \mid \tau_X(x_i) = \texttt{true}\} \cup \{\bar{x}_i \mid \tau_X(x_i) = \texttt{false}\} \\ &\quad \cup \{y_i \mid \tau_Y(y_i) = \texttt{true}\} \cup \{\bar{y}_i \mid \tau_Y(y_i) = \texttt{false}\}. \end{aligned}$$

It is easy to see that $S_{(\tau_X,\tau_Y)}$ is conflict-free in $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$. However, $S_{(\tau_X,\tau_Y)}$ cannot defend itself against all clause arguments $c_1, \dots, c_m$ in $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$, and therefore is not admissible: Since $\bar{\varphi}$ is in CNF and $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$, at least one clause in $\bar{\varphi}$ is unfulfilled. Let c_j be any such clause. Since the clauses of $\bar{\varphi}$ are disjunctions of literals, all literals in c_j are unfulfilled. The only arguments in $\mathscr{A}$ 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 $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$. 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 $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$, and we have $(\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle, S) \in$ PR-ATTINCPV.

For the other direction, assume that $(\varphi, X, Y) \notin \Sigma_2$SAT, i.e., $\forall \tau_X \exists \tau_Y : \bar{\varphi}[\tau_X, \tau_Y] = \texttt{true}$. Let τ_X be any assignment on X and let τ_Y be an assignment on Y that satisfies $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{true}$. Create the completion $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$ and the set $S_{(\tau_X,\tau_Y)}$ as before. Since $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{true}$, 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 $\langle \mathscr{A}, \mathscr{R}^{\tau_X} \rangle$, and since τ_X was generic, S is not preferred in any completion of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$, which proves $(\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle, S) \notin$ PR-ATTINCPV. ❑

The same hardness can be proven for the argument-incomplete model.

Theorem 18. PR-ARGINCPV *is* Σ_2^p*-hard.*

Proof. Again, we reduce from Σ_2SAT using a very similar construction. Given an instance (φ, X, Y) of Σ_2SAT, we create an instance $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S)$ of PR-ARGINCPV by setting $S = \emptyset$ and:

$$\mathscr{A} = \left\{ \begin{array}{ll} y_i, \bar{y}_i, & \text{for } y_i \in Y \\ \bar{x}_i, & \text{for } x_i \in X \\ c_i, & \text{for } c_i \text{ in } \bar{\varphi} \end{array} \right\}, \quad \mathscr{R} = \left\{ \begin{array}{ll} (\bar{y}_i, y_i), (y_i, \bar{y}_i), & \text{for } y_i \in Y \\ (x_i, \bar{x}_i), & \text{for } x_i \in X \\ (c_i, c_i), & \text{for } c_i \text{ in } \bar{\varphi} \\ (c_i, y_j), (c_i, \bar{y}_j), & \text{for } c_i \text{ in } \bar{\varphi}, y_j \in Y \\ (c_i, x_k), (c_i, \bar{x}_k), & \text{for } c_i \text{ in } \bar{\varphi}, x_k \in X \\ (y_j, c_i), & \text{if } y_j \text{ in } c_i \\ (\bar{y}_j, c_i), & \text{if } \neg y_j \text{ in } c_i \\ (x_k, c_i), & \text{if } x_k \text{ in } c_i \\ (\bar{x}_k, c_i), & \text{if } \neg x_k \text{ in } c_i \end{array} \right\}.$$

$$\mathscr{A}^? = \{\; x_i, \quad \text{for } x_i \in X \;\},$$

For an assignment τ_X on X, define the corresponding completion of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ by $\langle \mathscr{A}^{\tau_X}, \mathscr{R}|_{\mathscr{A}^{\tau_X}} \rangle$ with $\mathscr{A}^{\tau_X} = \mathscr{A} \cup \{x_i \in \mathscr{A}^? \mid \tau_X(x_i) = \texttt{true}\}$. This construction differs from that in the proof of Theorem 17 only in the implementation of the choice gadgets for the variables in X, which use possible arguments instead of possible attacks but which have the same effect: If, for $x_i \in X$, $\tau_X(x_i) = \texttt{true}$, then argument x_i is included in $\mathscr{A}^{\tau_X}$ and has an attack against argument $\bar{x}_i$ which S cannot defend, so x_i is a candidate for membership in admissible supersets of S and $\bar{x}_i$ is not. If $\tau_X(x_i) = \texttt{false}$, then x_i is excluded and does not attack $\bar{x}_i$, so $\bar{x}_i$ could be in admissible supersets of S. The remainder of the proof is completely analogous. ❑

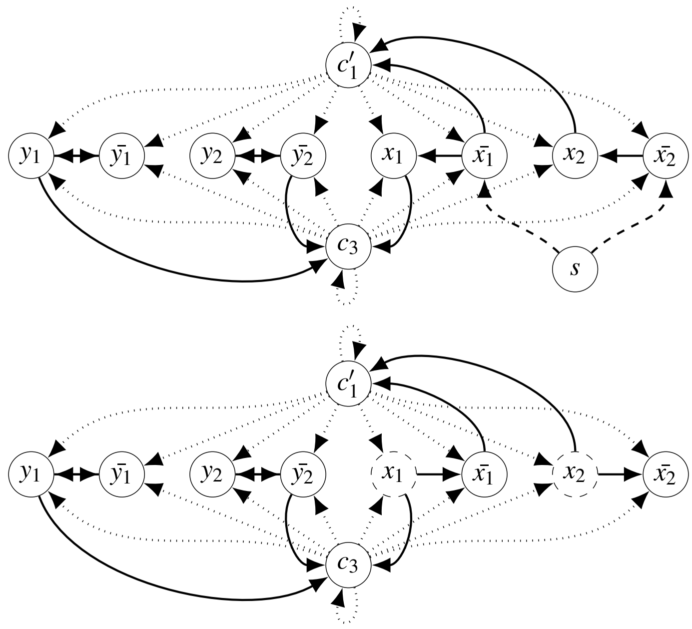

Figure 5: Yes-instances: Graph representations of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ (top) and $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ (bottom) created from clauses $c'_1 = (\neg x_1 \vee x_2)$ and $c_2 = (x_1 \vee y_1 \vee \neg y_2)$. Dashed attacks or arguments indicate uncertainty and attacks by clause arguments are displayed as dotted arcs to facilitate readability.

Example 19. Consider a Σ_2SAT 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 $(\varphi, X, Y) \notin \Sigma_2$SAT, because for all assignments τ_X on X and the assignment τ_Y with $\tau_Y(y_1) = \texttt{false}$, $\tau_Y(y_2) = \texttt{false}$ we have $\varphi[\tau_X, \tau_Y] = \texttt{false}$, or, equivalently, $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{true}$.

Figure 4 shows the graph representations of the incomplete argumentation frameworks in the instances $(\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle, \{s\})$ and $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, \emptyset)$ that are created from (φ, X, Y) according to the constructions in the proofs of Theorems 17 and 18. 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 $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$, while the set $\{\bar{y_1}, \bar{y_2}\}$ is an admissible superset of $\emptyset$ in all completions of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$.

To create a yes-instance, we slightly modify this Σ_2SAT instance by setting $\varphi' = (x_1 \wedge \neg x_2) \vee (\neg x_1 \wedge \neg y_1 \wedge y_2)$, i.e., $\neg y_1$ is omitted in the first clause. We now have $\bar{\varphi'} = \neg \varphi' = c'_1 \wedge c_2$, where $c'_1 = (\neg x_1 \vee x_2)$, and $c_2 = (x_1 \vee y_1 \vee \neg y_2)$ is unchanged. (φ', X, Y) is a yes-instance of Σ_2SAT, because for the assignment τ_X on X with $\tau_X(x_1) = \texttt{true}$, $\tau_X(x_2) = \texttt{false}$ and for all assignments τ_Y on Y, we have $\varphi[\tau_X, \tau_Y] = \texttt{true}$, or, equivalently, $\bar{\varphi}[\tau_X, \tau_Y] = \texttt{false}$.

Figure 5 shows the graph representations of the incomplete argumentation frameworks created from this modified Σ_2SAT instance. Both are yes-instances for PR-ATTINCPV and PR-ARGINCPV, respectively. The completions that correspond to the assignment τ_X as defined above include the possible attack $(s, \bar{x_1})$ (respectively, the possible argument x_1) and exclude the possible attack $(s, \bar{x_2})$ (respectively, the possible argument x_2). In these completions, there are no admissible supersets of S that attack c'_1, so S is preferred.

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

Corollary 20. PR-INCPV *is* Σ_2^p*-hard.*

5 Conclusion and Future Work

We extended prior research for three specific models of incompleteness in argumentation frameworks, i.e., *attack incompleteness* alone, *argument incompleteness* alone, and the combination of these two models so as to provide a *general model of incompleteness*. We studied, with respect to six common semantics of argumentation frameworks, the computational complexity of the possible and necessary verification problems, and filled gaps that have been left open by prior work.

Table 1: Overview of complexity results for various semantics (first column) in the standard model (second column), in the attack-incomplete model (third and sixth column), in the argument-incomplete model (fourth and seventh column), and in the combined model (fifth and eighth column). Results marked by ♠ are due to Dung [19], by ♣ due to Dimopoulos and Torres [18], by ★ due to Coste-Marquis et al. [16], by ▲ due to Baumeister et al. [6], by ▼ due to Baumeister et al. [7], and by ♦ due to Baumeister et al. [8]. For a complexity class $\mathscr{C}$, $\mathscr{C}$-c. stands for $\mathscr{C}$-completeness.

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

Table 1 gives an overview of the complexity results for the verification problem in the standard model and in the three incompleteness models considered in this paper. The results show a pattern in how incompleteness 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 model of argument incompleteness (and, therefore, also in the general incompleteness model). In all other cases—in particular, for all variants of necessary verification—introducing incomplete information does not make the verification problem computationally harder. Note that each of our hardness results for verification problems carries over to any more general model; so our approach is potentially useful in other frameworks as well. We further note that the Σ_2^p-completeness results for possible verification in the preferred semantics are significantly more severe than the NP- or coNP-completeness results for possible verification in the other semantics entailing admissibility and for standard or necessary verification in the preferred semantics: While there are known methods to circumvent NP- or coNP-hardness in practice (e.g., by using fast SAT-solvers), no such methods are effective to tame Σ_2^p-hardness in practice (even though there are also QBF-solvers, these are much less efficient in general). A task for future work is to analyze the complexity of possible and necessary variants of other decision problems than verification, e.g., credulous or skeptical acceptance of individual arguments. Also, the range of classical semantics considered here could be extended by including other, more recently proposed semantics like the stage semantics [34], the semi-stable semantics [13], the ideal semantics [21], or the CF2 semantics [1].

Acknowledgements

This paper merges and extends two preliminary ADT'15 papers [6, 7], a COMSOC'16 workshop paper [8], and an AAAI'18 paper [10], and we thank the reviewers for their helpful comments. This work was supported in part by the NRW Ministry for Innovation, Science, and Research.

References

[1] P. Baroni, M. Giacomin, and G. Guida. SCC-recursiveness: A general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

[2] R. Baumann and G. Brewka. Expanding Argumentation Frameworks: Enforcing and Monotonicity Results. In *Proceedings of the 3rd International Conference on Computational Models of Argument*, pages 75–86. IOS Press, September 2010.

[3] D. Baumeister, M. Roos, and J. Rothe. Computational Complexity of Two Variants of the Possible Winner Problem. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, pages 853–860. IFAAMAS, May 2011.

[4] D. Baumeister, M. Roos, J. Rothe, L. Schend, and L. Xia. The Possible Winner Problem with Uncertain Weights. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 133–138. IOS Press, August 2012.

[5] D. Baumeister, G. Erdélyi, O. Erdélyi, and J. Rothe. Complexity of Manipulation and Bribery in Judgment Aggregation for Uniform Premise-Based Quota Rules. *Mathematical Social Sciences*, 76:19–30, 2015.

[6] D. Baumeister, D. Neugebauer, and J. Rothe. Verification in Attack-Incomplete Argumentation Frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 341–358. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*, September 2015.

[7] D. Baumeister, J. Rothe, and H. Schadrack. Verification in Argument-Incomplete Argumentation Frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 359–376. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*, September 2015.

[8] D. Baumeister, D. Neugebauer, J. Rothe, and H. Schadrack. Verification in Incomplete Argumentation Frameworks. In *Proceedings of the 6th International Workshop on Computational Social Choice*, 2016.

[9] D. Baumeister, S. Bouveret, J. Lang, N. Nguyen, T. Nguyen, J. Rothe, and A. Saffidine. Positional Scoring-Based Allocation of Indivisible Goods. *Journal of Autonomous Agents and Multi-Agent Systems*, 31(3):628–655, 2017.

[10] D. Baumeister, D. Neugebauer, J. Rothe, and H. Schadrack. Complexity of Verification in Incomplete Argumentation Frameworks. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. AAAI Press, February 2018.

[11] G. Boella, S. Kaci, and L. van der Torre. 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*, pages 107–118. Springer-Verlag *Lecture Notes in Artificial Intelligence #5590*, July 2009.

[12] S. Bouveret, U. Endriss, and J. Lang. Fair Division under Ordinal Preferences: Computing Envy-Free Allocations of Indivisible Goods. In *Proceedings of the 19th European Conference on Artificial Intelligence*, pages 387–392. IOS Press, August 2010.

[13] M. Caminada. Semi-Stable Semantics. In *Proceedings of the 1st International Conference on Computational Models of Argument*, pages 121–130. IOS Press, September 2006.

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

[15] Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New Candidates Welcome! Possible Winners with respect to the Addition of New Candidates. *Mathematical Social Sciences*, 64 (1):74–88, 2012.

[16] S. Coste-Marquis, C. Devred, S. Konieczny, M. Lagasquie-Schiex, and P. Marquis. On the Merging of Dung's Argumentation Systems. *Artificial Intelligence*, 171(10):730–753, 2007.

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

[18] Y. Dimopoulos and A. Torres. Graph Theoretical Structures In Logic Programs and Default Theories. *Theoretical Computer Science*, 170(1):209–244, 1996.

[19] P. Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and *n*-Person Games. *Artificial Intelligence*, 77(2):321–357, 1995.

[20] P. Dung and P. Thang. Towards (Probabilistic) Argumentation for Jury-based Dispute Resolution. In *Proceedings of the 3rd International Conference on Computational Models of Argument*, pages 171–182. IOS Press, September 2010.

[21] P. Dung, P. Mancarella, and F. Toni. A Dialectic Procedure for Sceptical, Assumption-Based Argumentation. In *Proceedings of the 1st International Conference on Computational Models of Argument*, pages 145–156. IOS Press, September 2006.

[22] P. Dunne and M. Wooldridge. Complexity of Abstract Argumentation. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*, chapter 5, pages 85–104. Springer, 2009.

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

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

[25] K. Konczak and J. Lang. Voting Procedures with Incomplete Preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, pages 124–129, July/August 2005.

[26] J. Lang, A. Rey, J. Rothe, H. Schadrack, and L. Schend. Representing and Solving Hedonic Games with Ordinal Preferences and Thresholds. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 1229–1237. IFAAMAS, May 2015.

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

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

[29] M. Maher. Resistance to Corruption of Strategic Argumentation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 1030–1036. AAAI Press, February 2016.

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

[31] T. Rienstra. Towards a Probabilistic Dung-Style Argumentation System. In *Proceedings of the 1st International Conference on Agreement Technologies*, pages 138–152. CEUR Workshop Proceedings, October 2012.

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

[33] L. Stockmeyer. The Polynomial-Time Hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[34] Bart Verheij. Two Approaches to Dialectical Argumentation: Admissible sets and argumentation stages. In *Proceedings of the 8th Dutch Conference on Artificial Intelligence*, pages 357–368, 1996.

[35] J. Wallner, A. Niskanen, and M. Järvisalo. Complexity Results and Algorithms for Extension Enforcement in Abstract Argumentation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 1088–1094. AAAI Press, 2016.

[36] L. Xia and V. Conitzer. Determining Possible and Necessary Winners Given Partial Orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.

Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: {baumeister,neugebauer,rothe,schadrack}@cs.uni-duesseldorf.de

A Deferred Proof of Lemma 14

Proof of Lemma 14. If $S|_{\mathscr{A}_S^{ungr}}$ 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 $S|_{\mathscr{A}_S^{ungr}}$ be the grounded extension of IAF_S^{ungr}. We prove that, then, S is necessarily grounded in IAF.

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

Now, let $IAF^* = \langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ be any completion of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ (different from the ungrounded completion) and let G^* be its grounded extension. Since we know by assumption that $S|_{\mathscr{A}^*}$ is complete in IAF^*, with the fact (proven by Dung [19]) that the grounded extension is contained in all complete extensions of the same argumentation framework, we can conclude that $G^* \subseteq S|_{\mathscr{A}^*}$.

However, we also have $S|_{\mathscr{A}^*} \subseteq G^*$: Since $S|_{\mathscr{A}_S^{ungr}}$ contains only definite arguments, these must be in G^*, too. Now assume that $S|_{\mathscr{A}^*} \not\subseteq G^*$. Then there is a possible (nondefinite) argument $a \in (S|_{\mathscr{A}^*} \setminus G^*)$. We know that a is not included in the ungrounded completion. We also know that a

is not defended by G^* in IAF^*, because otherwise it would need to be included in the grounded set G^*. Also, since $S|_{\mathscr{A}_S^{ungr}} \subseteq G^*$, a is not defended by $S|_{\mathscr{A}_S^{ungr}}$ 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 $S|_{\mathscr{A}_S^{ungr}}$) in IAF^*. Since the ungrounded completion includes all arguments that are not in S, b is also included in $\mathscr{A}_S^{ungr}$. 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 $S|_{\mathscr{A}_S^{ungr}}$ in the ungrounded completion, which, by its construction, would mean that a would be included in $\mathscr{A}_S^{ungr}$ (a could only be excluded in Step 3 if it is defended by a subset of $S|_{\mathscr{A}_S^{ungr}}$, 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 $S|_{\mathscr{A}^*} \subseteq G^*$ and, in total, $S|_{\mathscr{A}^*} = G^*$. So, $S|_{\mathscr{A}^*}$ is grounded in IAF^* and, since IAF^* was kept generic, S is necessarily grounded in IAF. ❑ Lemma 14