Extension-based Semantics for Incomplete Argumentation Frameworks: Properties, Complexity and Algorithms

Jean-Guy Mailly
Université Paris Cité, LIPADE, F-75006 Paris, France

Abstract

Incomplete Argumentation Frameworks (IAFs) have been defined to incorporate some qualitative uncertainty in abstract argumentation: information such as "I am not sure whether this argument exists" or "I am not sure whether this argument attacks that one" can be expressed. Reasoning with IAFs is classically based on a set of completions, *i.e.* standard argumentation frameworks that represent the possible worlds encoded in the IAF. The number of these completions may be exponential with respect to the number of arguments in the IAF. This leads, in some cases, to an increase of the complexity of reasoning, compared to the complexity of standard AFs. In this paper, we follow an approach that was initiated for Partial AFs (a subclass of IAFs), which consists in defining new forms of conflict-freeness and defense, the properties that underly the definition of Dung's semantics for AFs. We generalize these semantics from PAFs to IAFs. We show that, among three possible types of admissibility, only two of them satisfy some desirable properties. We use them to define two new families of extension-based semantics. We study the properties of these semantics, and in particular we show that their complexity remains the same as in the case of Dung's AFs. Finally, we propose a logical encoding of these semantics, and we show experimentally that this encoding can be used efficiently to reason with IAFs, thanks to the power of modern SAT solvers.

1 Introduction

Abstract argumentation has been a major subfield of Knowledge Representation and Reasoning since the seminal paper by Dung [18]. However, although it is very appealing, Dung's framework is limited in the kind of information that can be modeled: only (abstract) arguments and attacks between them. For this reason, many generalizations of this framework have been proposed, introducing the notion of support between arguments [3], weighted attacks [19] or weighted arguments [44], preferences between arguments [2], and so on.

Among these generalizations of Dung's framework, a very natural research direction is the introduction of uncertainty in the model. Indeed, uncertainty is omnipresent in real world, and must be taken into account in the modeling of agents that reason about their environment or about other agents. Moreover, when arguments are generated from natural language processing [35], the nuances that exist in natural language are likely to be sources of uncertainty [6] that should appear in the formal model. Two directions have been followed for integrating uncertainty in abstract argumentation: quantitative representation of uncertainties (*e.g.* probabilities [36, 30]) and qualitative ones [16, 14, 7]. While quantitative representation of uncertainty is valuable when it is available, allowing fine grained reasoning about uncertainty, it may not be available in many realistic cases. For instance, in a debate, an agent can be uncertain that her opponent will use a given argument or not, without having a quantitative measure of this uncertainty. The study of qualitative models of uncertainty is thus of utter importance for the design of AI systems.

In this paper, we follow this direction. Qualitative uncertainty in abstract argumentation was originally studied in a context of Argumentation Framework (AF) merging [16]: Partial Argumentation Frameworks (PAFs) are AFs with possible ignorance about the existence of some attacks, initially used as a tool during some step of the merging process. Semantics dedicated to these PAFs were then defined in [14]. However, most of the work in this field focuses on a generalization of PAFs, namely Incomplete AFs (IAFs), where uncertainty concerns both the arguments and the attacks, and reasoning is based on *completions*. A completion is an argumentation framework that represents one of the (uncertain) scenarios encoded in the IAF. Classical reasoning tasks are then adapted in two versions: the possible view (is some property true for some completion?) and the necessary view (is some property true for each completion?). However, the number of completions is (in the worst case) exponential in the number of arguments. This means that various reasoning problems are harder for IAFs than their counterpart for standard AFs [9, 24, 7].

In this paper, we follow the approach initiated by [14]: we define new forms of conflict-freeness and defense based on the different types of information in an IAF. The combination of a notion of conflict-freeness and a notion of defense yields a notion of admissibility; we show that among the three possible variants of admissibility, only two of them satisfy some desirable property, namely Dung's Fundamental Lemma (adapted to IAFs). This lemma states, in classical AFs, that an admissible set remains admissible if an argument defended by it is added to the set. From the two "fundamental" notions of admissibility for IAFs (that we call weak and strong admissibility), we define (weak and strong) variants of the classical complete, preferred and stable semantics. We study some properties of these semantics, and we show that their complexity remains the same as in the standard AF case. Finally, we propose logical encodings of these semantics, in the same vein as [11]. We describe an implementation of our SAT-based approach for reasoning with the new semantics, and we empirically show that it scales up well.

This article is an extended version of a preliminary conference paper [39].

The new material included in this version is this one:

- we provide additional background notions, on propositional logic and computational complexity: Section 2.1 and Section 2.2,
- we correct an error included in the previous publication, proving that weak stable extensions are not weak admissible sets (but strong stable extensions are indeed strong admissible sets): Section 3.3;
- we provide new results on the relations between weak and strong variants of the semantics: Section 3.4);
- we study the complexity of additional decision problems (extension existence and non-emptiness) for all the semantics studied in the paper: Section 4.1.3 and Section 4.1.5;
- we describe an implementation of the SAT-based approach, and an experimental evaluation thereof: Section 5.

The rest of the paper is organized as follows. Section 2 describes background notions on propositional logic, computational complexity and abstract argumentation. In Section 3, we define our new semantics and study some of their properties, in particular the satisfaction of the Fundamental Lemma, and some inclusion relations between them. In Section 4, we show that the complexity remains the same as in the standard AF case,[1] and we provide a logical encoding for our semantics. Section 5 describes our implementation of the logical encoding defined in the previous section, and an experimental evaluation thereof shows that it scales up well. Finally, Section 6 describes some related work, and Section 7 concludes the paper.

2 Background

2.1 Propositional Logic and Boolean Satisfiability

We first recall some basic notions of classical logic, that will be useful in Section 4.2. We consider propositional formulas built on a set of Boolean variables V, *i.e.* each variable can be assigned a value in $\mathbb{B} = \{0, 1\}$ (where 0 is interpreted as *false*, and 1 as *true*). A well-formed formula is:

- $\phi = x$, for any $x \in V$ (atomic formula),
- $\phi = \neg\psi$, for ψ a well-formed formula (negation),
- $\phi = \psi \vee \psi'$, for ψ, ψ' two well-formed formulas (disjunction),
- $\phi = \psi \wedge \psi'$, for ψ, ψ' two well-formed formulas (conjunction).

[1] At the exception of skeptical acceptability under the complete semantics, for which we do not have a tight complexity result yet.

An interpretation is a mapping $\omega : V \to \mathbb{B}$, *i.e.* an assignment of a truth value to each variable. It can be extended to arbitrary formulas by the recursive mechanism:

- $\omega(\neg\phi) = 1 - \omega(\phi)$,
- $\omega(\phi \vee \psi) = \max(\omega(\phi), \omega(\psi))$,
- $\omega(\phi \wedge \psi) = \min(\omega(\phi), \omega(\psi))$.

An interpretation ω *satisfies* a formula ϕ if $\omega(\phi) = 1$. We also say that ω is a *model* of ϕ. We write $\mathsf{mod}(\phi)$ the set of models of ϕ. Finally, we define additional connectives as shortcuts for complex formulas:

- the material implication is $\phi \to \psi$, with $\mathsf{mod}(\phi \to \psi) = \mathsf{mod}(\neg\phi \vee \psi)$,
- the equivalence is $\phi \leftrightarrow \psi$, with $\mathsf{mod}(\phi \leftrightarrow \psi) = \mathsf{mod}((\phi \to \psi) \wedge (\psi \to \phi))$.

The *Boolean satisfiability problem* (SAT) consists in determining, given a propositional formula, whether it possesses at least one model. Although it is theoretically hard to solve in general (NP-complete [15], see Section 2.2), modern SAT solvers allow to solve it for many instances, including large ones [12]. This makes reductions to SAT a good method for solving many hard problems without developing specific algorithms for these problems. Notice that SAT solvers usually take as input Conjunctive Normal Form formulas (CNF), *i.e.* conjunction of clauses, where each clause is a disjunction of literals, and a literal is either an atomic formula, or the negation of an atomic formula. This is not a problem in practice, since any propositional formula can be translated into an equivalent CNF formula in polynomial time (modulo the addition of variables) [45].

2.2 Computational Complexity

We present now the basic notions of computational complexity that are used in the rest of this article. We focus on *decision problems*, *i.e.* questions that can be answered by "YES" or "NO". The goal is to determine how hard it is to solve such problems, with respect to the main resources required for computing a solution to these problems: time and space. To do so, we use the notion of *complexity class*, that are sets of problems sharing similar properties (*e.g.* being "easy" or "hard" to solve).

Roughly speaking, a problem is considered to be "easy" to solve (*tractable*) when there exists a deterministic algorithm that solves it in polynomial time with respect to the size of the problem instance, *i.e.* $\mathcal{O}(n^k)$ computation steps where n is the size of the instance, and $k \in \mathbb{N}$ is a fixed constant. These problems are gathered in the complexity class P. Among polynomial problems, we can identify space-logarithmic ones, that are the decision problems solved by a deterministic algorithm using a memory space logarithmic in the size of the input (besides the size of the input itself, naturally). This means that running the algorithm requires $\mathcal{O}(\log(n))$ memory units if the problem instance needs n memory units to be represented. These problems form the class L, which is a subset of P.

Among "hard" problems (*intractable*), special attention has been paid to NP, the set of problems that can be solved in polynomial time by a *non-deterministic* algorithm. A classical approach to identify a problem in NP is the following generic non-deterministic algorithm: given an instance I of the problem,

1. guess a potential proof p that I is a "YES" instance,
2. check (with a deterministic polynomial algorithm) that p is actually a proof that I is a "YES" instance.

The first step is called a non-deterministic guess. From this approach, NP is sometimes characterized as the set of problems for which it is hard to find a solution, but easy to verify a solution (step 2 of the algorithm).

If the second step does not use a deterministic polynomial algorithm, but a NP algorithm instead, it defines another class Σ_2^{P} (sometimes written $\mathsf{NP}^{\mathsf{NP}}$), the set of decision problems that can be solved in polynomial time by a non-deterministic algorithm with access to a NP oracle (*i.e.* a black box able to solve a problem from the class NP).

The complement of a complexity class $\mathcal{C}$ is $\overline{\mathcal{C}} = \{\overline{\mathcal{P}} \mid \mathcal{P} \in \mathcal{C}\}$, where $\overline{\mathcal{P}}$ is the complement problem of $\mathcal{P}$, *i.e.* the decision problem built on the same set of instances as $\mathcal{P}$, such that i is a "YES" instance of $\mathcal{P}$ if and only if it is a "NO" instance of $\overline{\mathcal{P}}$. Complement classes that will be used in the rest of this article are $\mathsf{coNP} = \overline{\mathsf{NP}}$ and $\Pi_2^{\mathsf{P}} = \overline{\Sigma_2^{\mathsf{P}}}$.

Decision problems in these complexity classes can be compared thanks to the notion of polynomial time reduction, *i.e.* a function f that takes as input instances of a problem $\mathcal{P}$, and outputs instances of a problem $\mathcal{P}'$, such that i is a "YES" instance of $\mathcal{P}$ if and only if $f(i)$ is a "YES" of $\mathcal{P}'$, and f is computable in polynomial time with respect to the size of i. In this case, we write $\mathcal{P} \leq_f^{\mathsf{P}} \mathcal{P}'$, which means that $\mathcal{P}'$ is at least as hard as $\mathcal{P}$. This notion is used to define the concept of $\mathcal{C}$-hardness: a problem $\mathcal{P}'$ is $\mathcal{C}$-hard if for any $\mathcal{P} \in \mathcal{C}$, $\mathcal{P} \leq_f^{\mathsf{P}} \mathcal{P}'$.[2] A problem which is $\mathcal{C}$-hard and belongs to $\mathcal{C}$ is called $\mathcal{C}$-complete, which means that it is one of the hardest problems in $\mathcal{C}$.

The classes mentioned here are part of the *polynomial hierarchy*, a family of complexity classes recursively defined from P, NP, and coNP, using the concept of oracles. Several inclusion relations exist between these classes, depicted by Figure 1.[3] Whether these inclusions are strict is still an open question; if they were not strict then we would say that the polynomial hierarchy *collapses*. However, the contrary is usually assumed, it means (for instance) that a $\Sigma_2\mathsf{P}$-complete problem is considered strictly harder than a NP-complete problem, which is in turn supposed to be strictly harder than a polynomial problem.

[2] If $\mathcal{C} = \mathsf{P}$, an additional constraint must be fulfilled, namely f must be computable using logarithmic space with respect to the size of i.

[3] The classes Δ_i^{P} correspond to problems that can be solved by using a polynomial number

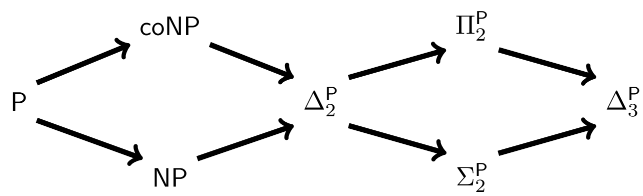

Figure 1: Illustration of the Polynomial Hierarchy first classes. $\mathcal{C}_1 \to \mathcal{C}_2$ is used to represent the inclusion of the complexity class $\mathcal{C}_1$ in the class $\mathcal{C}_2$.

For a more detailed overview of computational complexity, we refer the reader to *e.g.* [4].

2.3 Abstract Argumentation Frameworks

Abstract argumentation is the study of relations between abstract pieces of information called *arguments*; the internal nature of arguments, as well as their origin, is considered as irrelevant. Only the interactions between arguments are considered in order to determine which arguments are acceptable or not. The most classical type of relationship is the so-called *attack* relation, that expresses a contradiction between arguments. An attack is generally directed from one argument to another one, meaning that the first one somehow *defeats* the second one. The seminal paper [18] has launched the strong interest for abstract argumentation in the last 25 years. In this section, we formally introduce this abstract framework and how it is used for reasoning.

We suppose the existence of a finite set of arguments $\mathbf{A}$.

Definition 1 (Argumentation Framework)**.** *An argumentation framework (AF) is a pair $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ with $\mathcal{A} \subseteq \mathbf{A}$ the set of* arguments *and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ the set of* attacks*.*

For $a, b \in \mathcal{A}$, we say that *a attacks b* if $(a, b) \in \mathcal{R}$. If b attacks some $c \in \mathcal{A}$, then *a defends c* against b. Similarly, a set $S \subseteq \mathcal{A}$ attacks (respectively defends) an argument b if there is some $a \in S$ that attacks (respectively defends) b.

Example 1. *Figure 2 depicts an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, with $\mathcal{A} = \{a, b, c\}$ (*i.e.* the nodes of the graph) and $\mathcal{R} = \{(b, a), (b, c), (c, b)\}$ (*i.e.* the edges of the graph).*

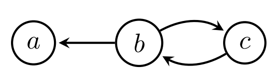

Figure 2: An Example of AF $\mathcal{F}$

The acceptability of arguments is classically evaluated through the concept of *extensions*, *i.e.* sets of arguments that are jointly acceptable. This form of

of calls to a Σ_i^{P} oracle. They are only shown for describing the position of Σ_2^P and Π_2^{P} in the hierarchy, but they not used in the rest of the article.

joint acceptance can be interpreted as defining a coherent point of view about the argumentative scenario that is represented by the AF. Different semantics have been defined, that yield different sets of extensions. The usual semantics are based on two main principles: conflict-freeness and admissibility.

Definition 2 (Conflict-freeness and Admissibility)**.** *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the set $S \subseteq \mathcal{A}$ is*

- conflict-free *iff* $\forall a, b \in S$, $(a, b) \notin \mathcal{R}$*;*
- admissible *iff it is conflict-free and* $\forall a \in S$, $\forall b \in \mathcal{A}$ *s.t.* $(b, a) \in \mathcal{R}$, $\exists c \in S$ *s.t.* $(c, b) \in \mathcal{R}$*.*

The meaning of conflict-freeness is quite easy to understand: we do not want to accept together arguments that are conflicting. Admissibility corresponds to a notion of "self-defense": a (conflict-free) set of arguments must be able to defend itself against external attacks in order to be considered as a valid point of view. We use $\mathsf{cf}(\mathcal{F})$ (respectively $\mathsf{ad}(\mathcal{F})$) to denote the set of conflict-free (respectively admissible) sets of an AF $\mathcal{F}$.

These principles are usually considered to be too weak to define semantics, but the classical semantics are based on them.[4] We recall now the definition of these semantics:

Definition 3 (Admissibility-based Semantics)**.** *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the admissible set $S \subseteq \mathcal{A}$ is*

- *a* complete *extension iff S contains all the arguments that it defends;*
- *a* preferred *extension iff S is a $\subseteq$-maximal admissible set;*
- *a* grounded *extension iff S is a $\subseteq$-minimal complete extension.*

A fourth semantics is defined by Dung, that does not directly rely on the notion of admissibility:

Definition 4 (Stable Semantics)**.** *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the conflict-free set $S \subseteq \mathcal{A}$ is a* stable extension *iff $\forall a \in \mathcal{A} \setminus S$, S attacks a.*

We use $\mathsf{co}(\mathcal{F})$, $\mathsf{pr}(\mathcal{F})$, $\mathsf{gr}(\mathcal{F})$ and $\mathsf{st}(\mathcal{F})$ for the sets of (respectively) complete, preferred, grounded and stable extensions. Among their basic properties:

- for any AF $\mathcal{F}$, $|\sigma(\mathcal{F})| \geq 1$ for $\sigma \in \{\mathsf{co}, \mathsf{pr}, \mathsf{gr}\}$;
- for any AF $\mathcal{F}$, $|\mathsf{gr}(\mathcal{F})| = 1$;
- for any AF $\mathcal{F}$, $\mathsf{st}(\mathcal{F}) \subseteq \mathsf{pr}(\mathcal{F}) \subseteq \mathsf{co}(\mathcal{F})$.

The last point implies that stable extensions are admissible sets as well, even if they are not explicitly defined through admissibility.

[4] However, let us notice that we will sometimes include them in the family of studied semantics, for homogeneity of the presentation, *e.g.* in the complexity results (see Section 4.1).

Example 2. *Considering again $\mathcal{F}$ from Example 1; its extensions for the four semantics defined previously are given in Table 1 (second column).*

For further details about these semantics, as well as other semantics that have been defined subsequently, we refer the reader to [18, 5].

Given an argumentation framework and a semantics, classical reasoning tasks include the verification that a given set of arguments is an extension, and that a given argument is credulously or skeptically acceptable, *i.e.* belongs to some or each extension. Formally:

σ-Ver Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $S \subseteq \mathcal{A}$, is S a σ-extension of $\mathcal{F}$?

σ-Cred Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $a \in \mathcal{A}$, does a belong to some σ-extension of $\mathcal{F}$?

σ-Skep Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $a \in \mathcal{A}$, does a belong to each σ-extension of $\mathcal{F}$?

We use $\mathsf{Cred}_\sigma(\mathcal{F})$ (respectively $\mathsf{Skep}_\sigma(\mathcal{F})$) to denote the set of credulously (respectively skeptically) accepted arguments of $\mathcal{F}$, *i.e.* those for which the answer to σ-Cred (respectively σ-Skep) is "YES".

Example 3. *The credulously and skeptically accepted arguments in $\mathcal{F}$ from Example 1 are given in Table 1 (third and fourth columns).*

Semantics σ	$\sigma(\mathcal{F})$	$\mathsf{Cred}_\sigma(\mathcal{F})$	$\mathsf{Skep}_\sigma(\mathcal{F})$
gr	$\{\emptyset\}$	$\emptyset$	$\emptyset$
st	$\{\{b\}, \{a,c\}\}$	$\{a,b,c\}$	$\emptyset$
co	$\{\emptyset, \{b\}, \{a,c\}\}$	$\{a,b,c\}$	$\emptyset$
pr	$\{\{b\}, \{a,c\}\}$	$\{a,b,c\}$	$\emptyset$

Table 1: Extensions and acceptable arguments of $\mathcal{F}$, for $\sigma \in \{\mathsf{gr}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$.

For most of the classical semantics, $\sigma(\mathcal{F}) \neq \emptyset$ holds for any $\mathcal{F}$. However, it is not the case for the stable semantics. This induces another interesting decision problem:

σ-Exist Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, is $\sigma(\mathcal{F}) \neq \emptyset$?

Finally, let us discuss the issue of (non-)emptiness in extension-based semantics. As said before, except for the stable semantics, most of the classical semantics always produce a non-empty set of extensions (*i.e.* $\sigma(\mathcal{F}) \neq \emptyset$ for any $\mathcal{F}$). However, there is no guarantee that there is a non-empty extension. On the contrary, the stable extensions may not exist, but if there are some then they are all non-empty. This conducts to the definition of the non-emptiness decision problem:

σ-NE Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, is there some $S \subseteq \mathcal{A}$ such that $S \neq \emptyset$ and $S \in \sigma(\mathcal{F})$?

The complexity of these problems for various semantics has been established, see *e.g.* [20] for an overview. The relevant results for this paper are summarized in Table 2.

Semantics σ	σ-Ver	σ-Cred	σ-Skep	σ-Exist	σ-NE
cf	in L	in L	trivial	trivial	in L
ad	in L	NP-c	trivial	trivial	NP-c
gr	P-c	P-c	P-c	trivial	in L
st	in L	NP-c	coNP-c	NP-c	NP-c
co	in L	NP-c	P-c	trivial	NP-c
pr	coNP-c	NP-c	Π_2^P-c	trivial	NP-c

Table 2: Complexity of σ-Ver, σ-Cred, σ-Skep, σ-Exist and σ-NE for $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{gr}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$. $\mathcal{C}$-c means $\mathcal{C}$-complete.

2.4 Qualitative Uncertainty in AFs

Now we present the existing models that incorporate qualitative uncertainty in abstract argumentation.

Incomplete Argumentation Frameworks

Definition 5 (Incomplete Argumentation Framework)**.** *An incomplete argumentation framework (IAF) is a tuple $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ where*

- $\mathcal{A} \subseteq \mathbf{A}$ *is the set of* certain arguments*;*
- $\mathcal{A}^? \subseteq \mathbf{A}$ *is the set of* uncertain arguments*;*
- $\mathcal{R} \subseteq (\mathcal{A} \cup \mathcal{A}^?) \times (\mathcal{A} \cup \mathcal{A}^?)$ *the set of* certain attacks*;*
- $\mathcal{R}^? \subseteq (\mathcal{A} \cup \mathcal{A}^?) \times (\mathcal{A} \cup \mathcal{A}^?)$ *the set of* uncertain attacks.

$\mathcal{A}$ *and* $\mathcal{A}^?$ *are disjoint sets of arguments, and* $\mathcal{R}$, $\mathcal{R}^?$ *are disjoint sets of attacks.*

Intuitively, $\mathcal{A}$ and $\mathcal{R}$ correspond, respectively, to arguments and attacks that certainly exist, while $\mathcal{A}^?$ and $\mathcal{R}^?$ are those that may (or may not) actually exist.

Example 4. *Figure 3 depicts an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ with $\mathcal{A} = \{a, b\}$ (plain nodes), $\mathcal{A}^? = \{c\}$ (square dashed node), $\mathcal{R} = \{(c, b)\}$ (plain edge) and $\mathcal{R}^? = \{(b, a)\}$ (dotted edge). It means that the arguments a and b certainly exist, and there is an uncertainty regarding the existence of the attack (b, a). Then, the argument c is uncertain, but if it exists then the attack (c, b) certainly exists as well.*

Reasoning with such IAFs is generally made through the notion of completion, *i.e.* a classical AF that represents a "possible world" with respect to the uncertain information encoded in the IAF:

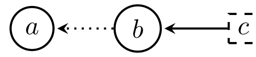

Figure 3: An Example of IAF $\mathcal{I}$

Definition 6 (Completion)**.** *Let $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an IAF. A completion of $\mathcal{I}$ is an AF $\mathcal{F}_c = \langle \mathcal{A}_c, \mathcal{R}_c \rangle$ such that*

- $\mathcal{A} \subseteq \mathcal{A}_c \subseteq \mathcal{A} \cup \mathcal{A}^?$*;*
- $\mathcal{R} \cap (\mathcal{A}_c \times \mathcal{A}_c) \subseteq \mathcal{R}_c \subseteq (\mathcal{R} \cup \mathcal{R}^?) \cap (\mathcal{A}_c \times \mathcal{A}_c)$*.*

Example 5. *Figure 4 depicts the completions of $\mathcal{I}$ from Example 4. $\mathcal{F}_1$ shows the situation where none of the uncertain elements actually exists, while $\mathcal{F}_4$ shows the opposite situation (all the uncertain elements appear). $\mathcal{F}_2$ and $\mathcal{F}_3$ shows the intermediate situations, where only one uncertain element (either the argument c, or the attack (b, a)) exists.*

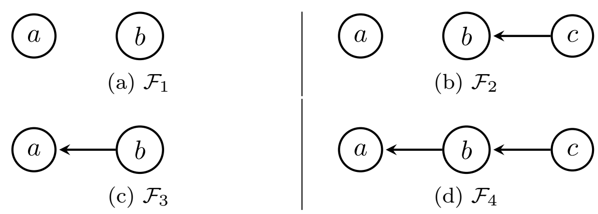

(a) $\mathcal{F}_1$ (b) $\mathcal{F}_2$ (c) $\mathcal{F}_3$ (d) $\mathcal{F}_4$

Figure 4: The Completions of $\mathcal{I}$

As seen with the previous example, the number of completions is generally exponential in the size of the IAF. More precisely, it is bounded by $2^{|\mathcal{A}^?|+|\mathcal{R}^?|}$.

Finally, reasoning tasks like credulous acceptance, skeptical acceptance or verification are defined with respect to some or each completion [9, 7]: each classical reasoning task has two variants, following the possible view (the property holds in some completion) and the necessary view (the property holds in each completion). These reasoning tasks are in many cases computationally harder than their counterpart for standard AFs (under the usual assumption that the polynomial hierarchy does not collapse) [9, 7]. This can be explained by the exponential number of completions.

Partial Argumentation Frameworks Partial Argumentation Frameworks were initially defined as a tool in a merging process [16]. They are tuples $\mathcal{P} = \langle A, R, I, N \rangle$ with three binary relations over the set of arguments A: R is the (certain) attack relation, I the ignorance relation, and N the (certain) non-attack relation. Since $N = (A \times A) \setminus (R \cup I)$, a PAF can be identified with only $\langle A, R, I \rangle$. Since the meaning of I is exactly the same as the meaning of $\mathcal{R}^?$, PAFs actually form a subclass of IAFs:[5] any PAF $\mathcal{P} = \langle A, R, I \rangle$ is equivalent to an IAF $\mathcal{I}_{\mathcal{P}} = \langle \mathcal{A}, \emptyset, \mathcal{R}, \mathcal{R}^? \rangle$ with $\mathcal{A} = A$, $\mathcal{R} = R$, $\mathcal{R}^? = I$.

[5]This subclass was studied under the name *Attack-Incomplete AFs* [8].

Extension-based semantics for PAFs have been defined in [14]. Intuitively, the idea consists in defining different forms of conflict-freeness and defense, and then combine them for defining three types of admissibility. From these new notions of admissibility, the authors define three variants of the preferred semantics, and study their properties. An interesting point is the fact that the complexity remains the same as in Dung's setting, contrary to the other reasoning methods for IAFs. These are the notions that are generalized from PAFs to IAFs in the next section.

3 Generalizing Extension-based Semantics from Partial to Incomplete AFs

In this section, we follow the same approach as [14] for defining semantics for IAFs. Instead of defining the extensions with respect to the set of completions of the IAF, we will generalize the basic concepts of conflict-freeness and defense to take into account the uncertainty in the IAF. Then, the usual admissibility-based semantics can be defined.

3.1 Conflict-free and Admissible Sets of IAFs

We follow two approaches for defining conflict-freeness and defense for IAFs:

- Optimistic view: we consider that only certain arguments and attacks are harmful, so keep the definition of conflict-freeness and defense as in Dung's frameworks;
- Pessimistic view: we consider that all attacks are harmful, and must be defended by certain arguments and attacks only.

By optimistic, we mean that the agent considers *e.g.* that $(a, b) \in \mathcal{R}^?$ does not make a a real "threat" against the acceptance of b. Roughly speaking, it means that the agent is tolerant to conflicts if they are uncertain. On the opposite, the pessimistic view means that the agent considers that all uncertain attacks against an argument are real threats against the acceptance of b, and that b must be defended by certain elements only in order to be accepted. Let us formally define the corresponding versions of conflict-freeness and defense.

Definition 7 (Weak and Strong Conflict-freeness)**.** *Let $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an IAF. The set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is*

- weakly conflict-free *iff $\forall a, b \in S \cap \mathcal{A}$, $(a, b) \notin \mathcal{R}$;*
- strongly conflict-free *iff $\forall a, b \in S$, $(a, b) \notin \mathcal{R} \cup \mathcal{R}^?$.*

We use $\mathsf{cf}_w(\mathcal{I})$ and $\mathsf{cf}_s(\mathcal{I})$ to denote, respectively, the weakly and strongly conflict-free sets of an IAF $\mathcal{I}$.

Example 6. *Figure 5 depicts an IAF* $\mathcal{I}_2 = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$*, with* $\mathcal{A} = \{a, b, d, e\}$*,* $\mathcal{A}^? = \{c, f\}$*,* $\mathcal{R} = \{(c, b), (e, b), (e, f)\}$ *and* $\mathcal{R}^? = \{(b, a), (b, e), (d, e)\}$*. The set* $\{a, b, c\}$ *is weakly conflict-free: the attack from* b *to* a *does not violate the weak conflict-freeness since it is uncertain, and the attack from* c *to* b *does not violate it either because the attacker (*c*) is uncertain. It is not strongly conflict-free because of the same two attacks. An example of strongly conflict-free set is* $\{a, c, e\}$*.*

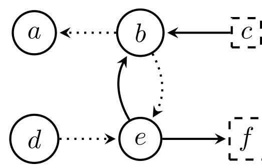

Figure 5: An Example of IAF $\mathcal{I}_2$

Strong conflict-freeness can be regarded as conflict-freeness applied on the "full" graph $\mathcal{F}_{full} = \langle \mathcal{A} \cup \mathcal{A}^?, \mathcal{R} \cup \mathcal{R}^? \rangle$, *i.e.* an AF made from the same arguments and attacks than the IAF, but without any uncertainty. However, weakly conflict-free sets do not correspond to the conflict-free sets of the "minimal" graph $\mathcal{F}_{min} = \langle \mathcal{A}, \mathcal{R} \cap (\mathcal{A} \times \mathcal{A}) \rangle$ (*i.e.* the AF obtained by simply ignoring the uncertain elements): see *e.g.* $\{a, b, c\}$ exhibited in Example 6, which is not a set of arguments in $\mathcal{F}_{min}$ (since $c \notin \mathcal{A}$).

Definition 8 (Weak and Strong Defense)**.** *Let* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *be an IAF. Given a set of arguments* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *and an argument* $a \in \mathcal{A} \cup \mathcal{A}^?$*,*

- S weakly defends a *iff* $\forall b \in \mathcal{A}$ *such that* $(b, a) \in \mathcal{R}$*,* $\exists c \in S \cap \mathcal{A}$ *s.t.* $(c, b) \in \mathcal{R}$*;*
- S strongly defends a *iff* $\forall b \in \mathcal{A} \cup \mathcal{A}^?$ *such that* $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$*,* $\exists c \in S \cap \mathcal{A}$ *s.t.* $(c, b) \in \mathcal{R}$*.*

Example 7. *Considering again* $\mathcal{I}_2$ *from Example 6, we observe that* $S = \{a\}$ *weakly defends* a*, since there is no* $x \in \mathcal{A}$ *s.t.* $(x, a) \in \mathcal{R}$*. On the contrary,* a *is not strongly defended by* S*, because there is no argument in* $S \cap \mathcal{A}$ *that attacks* b*. But* $S' = \{a, e\}$ *strongly defends* a*:* $e \in S' \cap \mathcal{A}$ *(certainly) attacks* b*.*

We observe that in the case where $\mathcal{A}^? = \emptyset$, then weak conflict-freeness and defense correspond to the notions of $\mathcal{R}$-conflict-freeness and $\mathcal{R}$-acceptability defined in [14], while the strong versions correspond to $\mathcal{RI}$-conflict-freeness and $\mathcal{RI}$-acceptability. Thus, if $\mathcal{R}^? = \emptyset$ also holds, then both weak conflict-freeness and strong conflict-freeness coincide with the classical conflict-freeness [18], while both forms of defense defined here correspond with the classical defense.

For defining a notion of admissibility, we must combine conflict-freeness and defense. In theory, Definitions 7 and 8 induce four notions of admissibility. However, the following result shows that weak conflict-freeness and strong conflict-freeness induce the same notion of admissibility when combined with strong defense.

Proposition 1. *Let $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an IAF. Let $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ be a set of arguments such that S is weakly conflict-free and $\forall a \in S$, S strongly defends a. Then S is strongly conflict-free.*

The proof is similar to the proof of [14, Property 1].

Proof. Reasoning towards a contradiction, let us suppose that S is not strongly conflict-free, *i.e.* $\exists a, b \in S$ such that $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$. Then, since S strongly defends all its elements, in particular it strongly defends b, so $\exists c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$. Now there are two options: either $a \in \mathcal{A}$ or $a \in \mathcal{A}^?$. First assume $a \in \mathcal{A}$. In that case, there is a contradiction between the existence of the attack (c, a) and the weak conflict-freeness of S. Now, assume that $a \in \mathcal{A}^?$. Since it is assumed that S strongly defends all its elements, there must be some $d \in S \cap \mathcal{A}$ such that $(d, c) \in \mathcal{R}$ (*i.e.* a is strongly defended against its attacker c). Now, we have a certain attack (d, c) between two certain arguments $c, d \in S \cap \mathcal{A}$, which is in contradiction with the weak conflict-freeness of S.

So we can conclude that S is strongly conflict-free. $\square$

Now we define the three variants of admissibility.[6]

Definition 9 (Weak, Mixed and Strong Admissibility)**.** *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, a set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is*

- weakly admissible *iff S is weakly conflict-free and weakly defends all its elements;*
- mixedly admissible *iff S is strongly conflict-free and weakly defends all its elements;*
- strongly admissible *iff S is strongly conflict-free and strongly defends all its elements.*

The weakly (respectively mixedly, strongly) admissible sets of an IAF $\mathcal{I}$ are denoted by $\mathsf{ad}_w(\mathcal{I})$ (respectively $\mathsf{ad}_m(\mathcal{I})$, $\mathsf{ad}_s(\mathcal{I})$).

The definitions imply that $\mathsf{ad}_s(\mathcal{I}) \subseteq \mathsf{ad}_m(\mathcal{I}) \subseteq \mathsf{ad}_w(\mathcal{I})$, for any IAF $\mathcal{I}$. Also, as in the standard Dung's framework, every IAF has at least one admissible set, for all the variations of admissibility. Indeed, for any IAF $\mathcal{I}$, $\emptyset \in \mathsf{ad}_s(\mathcal{I})$. This fact will be useful later to guarantee the existence of extensions for the semantics based on admissibility.

Before going further with the definition of semantics based on these new notions of admissibility, we briefly discuss a property of classical semantics that we believe is important. It is called the *Fundamental Lemma* by Dung [18, Lemma 10]. This lemma states that if a set of arguments S is admissible, and

[6] The terminology "strong defense" and "strong admissibility" has been used with another meaning in [13], where it applies to classical AFs, not IAFs.

defends an argument a, then $S \cup \{a\}$ is admissible. Besides its technical interest for proving some further results, this lemma describes an intuitive property of argumentation in general: if a point of view (*i.e.* a set of arguments) is seen as valid, then it should be jointly acceptable with any argument that it successfully defends. We thus consider this property as necessary for defining reasonable semantics. With the following lemma, we determine which of the notions of admissibility given in Definition 9 satisfy a notion of "fundamentality" similar to Dung's lemma. More precisely, we show that only weak and strong admissibility are suitable for defining semantics.

Lemma 1 (Fundamental Lemma). *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, and $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ a weakly (respectively strongly) admissible set, if S weakly (respectively strongly) defends some $a \in \mathcal{A} \cup \mathcal{A}^?$, then $S \cup \{a\}$ is weakly (respectively strongly) admissible.*

Proof. We first consider weak admissibility. Let us prove that $S \cup \{a\}$ is weakly conflict-free. First of all, notice that if $a \in \mathcal{A}^?$ then the set $S \cup \{a\}$ is weakly conflict-free iff S is weakly conflict-free, since only certain attacks between certain arguments violate weak conflict-freeness. So in the rest of the reasoning we suppose that $a \in \mathcal{A}$. Towards a contradiction, suppose that $S \cup \{a\}$ is not weakly conflict-free. Then, $\exists b \in S \cap \mathcal{A}$ such that, either $(b, a) \in \mathcal{R}$ or $(a, b) \in \mathcal{R}$. In the former case, since S weakly defends a, then there must be a $c \in S \cap \mathcal{A}$ with $(c, b) \in \mathcal{R}$, which is impossible since S is weakly conflict-free. Hence the contradiction. In the latter case ($(a, b) \in \mathcal{R}$), since S is weakly admissible, it must defend b against a, and the same reasoning applies for concluding the impossibility. Thus $S \cup \{a\}$ is weakly conflict-free. The fact that $S \cup \{a\}$ weakly defends all its elements comes from the fact that S weakly defends all its elements, as well as a. So we conclude that $S \cup \{a\}$ is weakly admissible.

Now, consider S a strongly admissible set that strongly defends some $a \in \mathcal{A} \cup \mathcal{A}^?$. Suppose that $S \cup \{a\}$ is not strongly conflict-free. It means that some $b \in S$ is such that $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$ or $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$. In the first case, the fact that S strongly defends a (against b) means that some $c \in S \cap \mathcal{A}$ attacks b, which violates strong conflict-freeness of S. In the second case, since S strongly defends all its elements, there is a $c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$, which is impossible for similar reasons to the first case. Hence $S \cup \{a\}$ is strongly conflict-free. Finally, the fact that $S \cup \{a\}$ strongly defends all its elements follows the fact that S strongly defends all its elements and a. So we conclude that $S \cup \{a\}$ is strongly admissible. $\square$

On the contrary, mixed admissibility does not satisfy a property of fundamentality.

Proposition 2. *There is an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ and an argument $a \in \mathcal{A} \cup \mathcal{A}^?$ such that S is mixedly admissible, S weakly defends a, and $S \cup \{a\}$ is not mixedly admissible.*

Proof. The IAF given at Figure 6 provides an example. The set $S = \{b\}$ is mixedly admissible (it is strongly conflict-free, and it has no attacker). S

weakly defends a (since there is no $x \in \mathcal{A}$ such that $(x, a) \in \mathcal{R}$, there is actually no need to weakly defend a). But $S \cup \{a\}$ is not strongly conflict-free, hence not mixedly admissible. $\square$

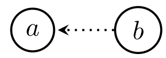

Figure 6: A Counter-Example about Fundamentality of Mixed Admissibility

Because of this reason, we do not consider mixed admissibility as suitable for defining semantics (*e.g.* mixed preferred or mixed complete semantics).

Example 8. *Based on Example 6 and 7, we observe that, in $\mathcal{I}_2$ from Figure 5, $\{a\}$ is weakly admissible but not strongly admissible. $\{a, e\}$ is not strongly admissible either, because it does not strongly defend e (against the uncertain attack (d, e)). The full sets of weakly and strongly admissible sets of $\mathcal{I}_2$ are given in Table 3.*

$x \in \{w, s\}$	w	s
$\mathsf{ad}_x(\mathcal{I}_2)$	$\emptyset$, $\{a\}$, $\{c\}$, $\{d\}$, $\{e\}$, $\{a, c\}$, $\{a, d\}$, $\{a, e\}$, $\{c, d\}$, $\{c, e\}$, $\{d, e\}$, $\{a, c, d\}$, $\{a, c, e\}$, $\{a, d, e\}$, $\{c, d, e\}$, $\{a, c, d, e\}$	$\emptyset$, $\{c\}$, $\{d\}$, $\{c, d\}$

Table 3: Weakly and Strongly Admissible Sets of $\mathcal{I}_2$.

3.2 Admissibility-based Semantics for IAFs

The classical definitions of Dung's semantics can be adapted to IAFs, based on the two different notions of admissibility identified as suitable in Lemma 1.

Definition 10 (Admissibility-based Semantics)**.** *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, a weakly (respectively strongly) admissible set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is*

- *a* weakly (respectively strongly) complete *extension iff S contains all the arguments that it weakly (respectively strongly) defends;*
- *a* weakly (respectively strongly) preferred *extension iff it is a $\subseteq$-maximal weakly (respectively strongly) admissible set.*

For $x \in \{w, s\}$ and $\sigma \in \{\mathsf{co}, \mathsf{pr}\}$, the set of x-σ extensions of an IAF $\mathcal{I}$ is denoted $\sigma_x(\mathcal{I})$. In the definition of the versions of complete semantics, the notion of defense used is the same as in the underlying notion of admissibility.

Example 9. *We continue Example 8. From the weakly and strongly admissible sets described in Table 3, we deduce $\mathsf{co}_w(\mathcal{I}_2) = \mathsf{pr}_w(\mathcal{I}_2) = \{\{a, c, d, e\}\}$, and $\mathsf{co}_s(\mathcal{I}_2) = \mathsf{pr}_s(\mathcal{I}_2) = \{\{c, d\}\}$.*

We observe some usual properties regarding these semantics.

Proposition 3. *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, and $x \in \{w, s\}$,*

- $\mathsf{pr}_x(\mathcal{I}) \neq \emptyset$*;*
- $\mathsf{pr}_x(\mathcal{I}) \subseteq \mathsf{co}_x(\mathcal{I})$*.*

Proof. The first item is a direct consequence of the fact that $\mathsf{ad}_x(\mathcal{I}) \neq \emptyset$, as seen previously. The existence of (finitely many) admissible sets implies the existence of $\subseteq$-maximal admissible sets.

Now, let S be a x-preferred extension of $\mathcal{I}$. Reasoning towards a contradiction, let us suppose that $S \notin \mathsf{co}_x(\mathcal{I})$. Since S is x-admissible, it means that S x-defends some argument a that it does not contain. According to Lemma 1, $S \cup \{a\}$ is x-admissible. This means that we have identified a proper superset of S which is x-admissible, thus S is not a $\subseteq$-maximal x-admissible set. This contradicts the fact that S is x-preferred. So we can conclude $S \in \mathsf{co}_x(\mathcal{I})$. □

3.3 Stable Semantics for IAFs

Now we focus on a counterpart of stable semantics for IAFs.

Definition 11 (Stable Semantics)**.** *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF,*

- *a weakly conflict-free set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is a* weakly stable *extension iff $\forall a \in \mathcal{A} \setminus S$, there is some $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$;*
- *a strongly conflict-free set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is a* strongly stable *extension iff $\forall a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$, there is some $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$.*

Weakly and strongly stable extensions of an IAF $\mathcal{I}$ are denoted by $\mathsf{st}_x(\mathcal{I})$, where $x \in \{w, s\}$.

Example 10. *Continuing Example 9, we observe that the weakly preferred extension $S = \{a, c, d, e\}$ is weakly stable as well: the argument $e \in S \cap \mathcal{A}$ (certainly) attacks all the arguments in $\mathcal{A} \setminus S = \{b\}$. It is not strongly stable, since it is not strongly conflict-free. The same applies to $\{a, d, e\}, \{a, d, e, f\}$ and $\{a, c, d, e, f\}$ which are also weakly conflict-free and certainly attack b.*

*On the contrary, the strongly preferred extension $S' = \{c, d\}$ is not strongly stable, since it does not attack all the arguments in $(\mathcal{A} \cup \mathcal{A}^?) \setminus S$ (*e.g. *a is not attacked by S').*

In Dung's framework, although admissibility is not directly involved in the definition of the stable semantics, any stable extension is actually an admissible set. Example 10 shows that weak stable extensions are not weakly admissible in general: $\{a, c, d, e, f\}$ is weakly conflict-free, but it does not weakly defend f, thus it is not weakly admissible. But we prove here that it is the case for strong stable semantics of IAFs.

Proposition 4 (Admissibility of Strong Stable Extensions)**.** *For any IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, $\mathsf{st}_s \subseteq \mathsf{ad}_s(\mathcal{I})$.*

Proof. Consider S a strongly stable extension of $\mathcal{I}$. Again, strong conflict-freeness is implied by the definition, so we just need to prove that S strongly defends all its elements. Consider any $a \in \mathcal{A} \cup \mathcal{A}^?$ such that $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$, for some $b \in S$. By definition of strongly stable extensions, there is some $c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$. Thus S strongly defends a, and then all its elements. We can conclude that it is strongly admissible. $\square$

Another classical result that still holds for the strong stable semantics semantics is the relationship between stable and preferred extensions.

Proposition 5 (Preferredness of Strong Stable Extensions)**.** *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$*,* $\mathsf{st}_s \subseteq \mathsf{pr}_s(\mathcal{I})$*.*

Proof. Consider $S \in \mathsf{st}_s(\mathcal{I})$. Proposition 4 implies the strong admissibility of S. Suppose the existence of $S' \in \mathsf{ad}_s(\mathcal{I})$ with $S \subset S'$. Take $a \in S' \setminus S$; the strong stability of S implies the existence of $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$, thus violating the strong admissibility of S'. We reach a contradiction, and conclude that S' does not exist, hence $S \in \mathsf{pr}_s(\mathcal{I})$. $\square$

Example 10 and Proposition 5 imply that $\mathsf{st}_s(\mathcal{I}_2) = \emptyset$. The non-existence of stable extensions in Dung's framework is one of the main differences between this semantics and the ones based on admissibility. We can simply show a similar example for the weakly stable semantics as well: the IAF $\mathcal{I} = \langle \{a\}, \emptyset, \{(a, a)\}, \emptyset \rangle$ has a single weakly conflict-free set (the empty set), which is not weakly stable.

3.4 Relations between Weak and Strong Semantics

From Definition 7, we can observe that strong conflict-freeness implies weak conflict-freeness. Formally,

Observation 1. *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$*,* $\mathsf{cf}_s(\mathcal{I}) \subseteq \mathsf{cf}_w(\mathcal{I})$*.*

In this part of the paper, we establish similar relationships between the weak and strong variants of other semantics. Indeed, the same observation can be made about the two concepts of defense.

Observation 2. *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$*,* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *and* $a \in \mathcal{A} \cup \mathcal{A}^?$*, if* S *strongly defends* a *then* S *weakly defends* a*.*

These two observations imply that strong σ-extensions are also weak σ-extensions for some of the semantics studied in this paper.

Proposition 6. *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *and* $\sigma \in \{\mathsf{ad}, \mathsf{st}\}$*,* $\sigma_s(\mathcal{I}) \subseteq \sigma_w(\mathcal{I})$*.*

Proof. For $\sigma = \mathsf{ad}$, the proof is obvious, following Observations 1 and 2.

Now, consider $E \in \mathsf{st}_s(\mathcal{I})$. From the definition, $E \in \mathsf{cf}_s(\mathcal{I}) \subseteq \mathsf{cf}_w(\mathcal{I})$. Then, we know that for each $a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$, some $b \in S \cap \mathcal{A}$ certainly attacks a. This implies that for each $a \in \mathcal{A} \setminus S$, some $b \in S \cap \mathcal{A}$ attacks a, and then S is weakly stable. $\square$

However, this result is not true for the complete and preferred semantics.

Example 11. *Consider $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ depicted at Figure 7. The strongly admissible sets of $\mathcal{I}$ are $\mathsf{ad}_s(\mathcal{I}) = \{\emptyset, \{b\}\}$, and the weakly admissible sets are $\mathsf{ad}_w(\mathcal{I}) = \{\emptyset, \{a\}, \{b\}, \{a, b\}\}$, hence $\mathsf{pr}_s(\mathcal{I}) = \{\{b\}\} \nsubseteq \mathsf{pr}_w(\mathcal{I}) = \{\{a, b\}\}$. The reasoning applies for the complete semantics as well, since here $\mathsf{co}_x(\mathcal{I}) = \mathsf{pr}_x(\mathcal{I})$, for $x \in \{s, w\}$.*

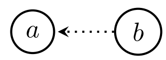

Figure 7: Counterexample for the relationship between strong and weak σ semantics, $\sigma \in \{\mathsf{co}, \mathsf{pr}\}$

However, Proposition 7 implies some relationship between strong complete (and preferred) extensions and weak complete (and preferred) extensions: indeed, each strong extension is included in a weak one.

Proposition 7. *For any IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $\sigma \in \{\mathsf{co}, \mathsf{pr}\}$, for each $S \in \sigma_s(\mathcal{I})$, there is a $S' \in \sigma_w(\mathcal{I})$ such that $S \subseteq S'$.*

Proof. Let S be a strong complete (or preferred) extension of $\mathcal{I}$. From the definition of the semantics, S is strongly admissible, then from Proposition 6 S is weakly admissible. This implies that S is included in some weakly preferred extension (which are the $\subseteq$-maximal weakly admissible sets), which is a particular weak complete extension (Proposition 3). Hence the result. $\square$

We summarize the inclusion schemes between the semantics in Table 4. For a cell with coordinates (σ_1, σ_2),

- ✓ means that, for any $\mathcal{I}$, $\sigma_1(\mathcal{I}) \subseteq \sigma_2(\mathcal{I})$,
- ✠ means that, for any $\mathcal{I}$, $\forall S \in \sigma_1(\mathcal{I})\ \exists S' \in \sigma_2(\mathcal{I})$ such that $S \subseteq S'$.

4 Computational Issues

4.1 Computational Complexity

In this section, we study the complexity of the variants of verification, existence, credulous acceptability, skeptical acceptability and non-emptiness for IAFs. Formally, for $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{co}, \mathsf{pr}, \mathsf{st}\}$ and $x \in \{w, s\}$:

σ_x-Ver Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $S \subseteq \mathcal{A}$, is $S \in \sigma_x(\mathcal{I})$?

σ_x-Exist Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, is $\sigma_x(\mathcal{I}) \neq \emptyset$?

σ_x-Cred Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $a \in \mathcal{A} \cup \mathcal{A}^?$, does a belong to some x-σ-extension of $\mathcal{I}$?

σ_1 \ σ_2	cf_w	cf_s	ad_w	ad_s	co_w	co_s	pr_w	pr_s	st_w	st_s
cf_w	✓									
cf_s	✓	✓								
ad_w	✓		✓							
ad_s	✓	✓	✓	✓						
co_w	✓		✓		✓					
co_s	✓	✓	✓	✓	$\maltese$	✓	$\maltese$			
pr_w	✓		✓		✓		✓			
pr_s	✓	✓	✓	✓	$\maltese$	✓	$\maltese$	✓		
st_w	✓								✓	
st_s	✓	✓	✓	✓	$\maltese$	✓	$\maltese$	✓	✓	✓

Table 4: Summary of the Inclusions between Semantics

σ_x-Skep Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $a \in \mathcal{A} \cup \mathcal{A}^?$, does a belong to each x-σ-extension of $\mathcal{I}$?

σ_x-NE Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, is there some $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ such that $S \neq \emptyset$ and $S \in \sigma_x(\mathcal{I})$?

4.1.1 Lower Bounds

We can prove that reasoning with our semantics for IAFs is (at least) as hard as reasoning with the corresponding semantics for AFs. This can be done by showing that any AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ can be transformed into an IAF $\mathcal{I}_\mathcal{F}$ that has the same extensions.

Definition 12 (IAF Associated with an AF)**.** *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the IAF associated with $\mathcal{F}$ is $\mathcal{I}_\mathcal{F} = \langle \mathcal{A}, \emptyset, \mathcal{R}, \emptyset \rangle$.*

Now we prove the correspondence of extensions, *i.e.* $\sigma(\mathcal{F}) = \sigma_w(\mathcal{I}_\mathcal{F}) = \sigma_s(\mathcal{I}_\mathcal{F})$, for any $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$.

Proposition 8 (Dung Compatibility)**.** *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$ and $x \in \{w, s\}$, $\sigma(\mathcal{F}) = \sigma_x(\mathcal{I}_\mathcal{F})$, where $\mathcal{I}_\mathcal{F}$ follows Definition 12.*

Proof. Observe that a set $S \subseteq \mathcal{A}$ is conflict-free (in $\mathcal{F}$) iff it is weakly and strongly conflict-free (in $\mathcal{I}_\mathcal{F}$). Then, a set $S \subseteq \mathcal{A}$ defends an argument $a \in \mathcal{A}$ against all it attackers (in $\mathcal{F}$) iff it weakly and strongly defends a against all its attackers (in $\mathcal{I}_\mathcal{F}$). These facts imply $\mathsf{ad}(\mathcal{F}) = \mathsf{ad}_w(\mathcal{I}_\mathcal{F}) = \mathsf{ad}_s(\mathcal{I}_\mathcal{F})$, which in turn imply the equivalence of complete and preferred extensions of $\mathcal{F}$ with the (weak and strong) complete and preferred extensions of $\mathcal{I}_\mathcal{F}$. Given $S \subseteq \mathcal{A}$, the equivalence between the conditions for S being stable in $\mathcal{F}$ and (weakly or strongly) stable in $\mathcal{I}_\mathcal{F}$ is straightforward. $\square$

This allows to prove that the complexity of reasoning with AFs provides a lower bound of the complexity of reasoning with IAFs.

Proposition 9. *Given* $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$*,* $x \in \{w, s\}$*, and* $\mathcal{P} \in \{\mathsf{Ver}, \mathsf{Exist}, \mathsf{Cred}, \mathsf{Skep}, \mathsf{NE}\}$*, if* σ*-*$\mathcal{P}$ *is* $\mathcal{C}$*-hard, then* σ_x*-*$\mathcal{P}$ *is* $\mathcal{C}$*-hard.*

Proof. Proposition 8 provides a polynomial-time and logarithmic-space reduction from σ-$\mathcal{P}$ to σ_x-$\mathcal{P}$. □

4.1.2 Upper Bounds for Extension Verification

Similarly to Dung's classical setting, most of the properties of extensions can be verified in polynomial time for our IAF semantics.

Lemma 2. *Given an IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *and a set of arguments* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$*, the following tasks are doable in polynomial time and logarithmic space:*

1. *check whether* S *is weakly (respectively strongly) conflict-free,*

2. *check whether* S *weakly (respectively strongly) defends some argument* $a \in \mathcal{A}$ *(respectively* $a \in \mathcal{A} \cup \mathcal{A}^?$*),*

3. *check whether each argument in* $\mathcal{A} \setminus S$ *(respectively* $(\mathcal{A} \cup \mathcal{A}^?) \setminus S$*) is attacked by an argument in* $S \cap \mathcal{A}$*.*

Proof. For item 1., weak (respectively strong) conflict-freeness is checked by iterating over $\{(a, b) \in S \times S\}$, and verifying whether $(a, b) \in \mathcal{R}$ (respectively $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$). There are $|S|^2$ such pairs (a, b), and verifying the membership to $\mathcal{R}$ (respectively $\mathcal{R} \cup \mathcal{R}^?$) is bounded by $|\mathcal{A} \cup \mathcal{A}^?|^2$ (*i.e.* the maximal number of possible attacks in an IAF).

For item 2., identifying the arguments $b \in \mathcal{A}$ (respectively $b \in \mathcal{A} \cup \mathcal{A}^?$) such that $(b, a) \in \mathcal{R}$ (respectively $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$) only requires to iterate over the arguments in $\mathcal{A}$ (respectively $\mathcal{A} \cup \mathcal{A}^?$), and then polynomially check the membership to $\mathcal{R}$ (respectively $\mathcal{R} \cup \mathcal{R}^?$). Then, for each of these attackers b, iterate over the arguments $c \in S \cap \mathcal{A}$ and check the membership of (c, b) to $\mathcal{R}$ (respectively $\mathcal{R} \cup \mathcal{R}^?$). All the iterations are polynomially bounded.

Finally, for item 3., enumerate all the pairs (a, b) such that $a \in S \cap \mathcal{A}$ and $b \in \mathcal{A} \setminus S$ (respectively $b \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$), and then check whether $(a, b) \in \mathcal{R}$. □

Combining these operations allows to check whether a set of arguments is an extension, for most of the semantics studied in this paper.

Proposition 10. *For* $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{co}, \mathsf{st}\}$ *and* $x \in \{w, s\}$*,* σ_x-Ver *is doable in polynomial time and logarithmic space.*

Proof. The result straightforwardly follows Lemma 2. □

Following Proposition 9, the verification of (weakly or strongly) preferred extensions is intractable (under the usual assumptions of complexity theory). The following results proves that it remains at the first level of the polynomial hierarchy, similarly to Dung's preferred semantics.

Proposition 11. *For* $x \in \{w, s\}$*,* pr_x-Ver *is in* coNP*.*

Proof. Given $S \subseteq \mathcal{A} \cup \mathcal{A}^?$, proving that S is not a weakly (respectively strongly) preferred extension is doable with the following non-deterministic polynomial algorithm:

1. Check whether S is weakly (respectively strongly) admissible. If not, then S is not weakly (respectively strongly) preferred.
2. Otherwise, guess a proper superset of S, *i.e.* $S \subset S' \subseteq \mathcal{A} \cup \mathcal{A}^?$. Verifying whether S' is a weakly (respectively strongly) admissible set is doable in polynomial time with a deterministic algorithm. If S' is weakly (respectively strongly) admissible, then S is not a weakly (respectively strong) preferred extension.

This algorithm proves that the complementary problem is in NP, thus we conclude that pr_x-$\mathsf{Ver} \in \mathsf{coNP}$ for $x \in \{w, s\}$. □

4.1.3 Upper Bounds for Existence

We recall that $\emptyset$ is weakly and strongly admissible (and naturally, weakly and strongly conflict-free as well) for any IAF. This implies that for any $\mathcal{I}$, $\mathsf{ad}_x(\mathcal{I}) \neq \emptyset$, for $x \in \{w, s\}$. The existence of some $\subseteq$-maximal elements in $\mathsf{ad}_x(\mathcal{I})$ is then guaranteed, *i.e.* $\mathsf{pr}_x(\mathcal{I}) \neq \emptyset$, and finally since $pr_x(\mathcal{I}) \subseteq \mathsf{co}_x(\mathcal{I})$, we obtain $\mathsf{co}_x(\mathcal{I}) \neq \emptyset$ as well. This means that our semantics have another common point with their counterpart in Dung's framework: all of them, except the stable semantics, induce a non-empty set of extensions for any IAF. We show that the question of existence for the stable semantics is NP-complete (NP-hardness follows Proposition 9, so we focus on NP-membership).

Proposition 12. *For $x \in \{w, s\}$, st_x-Exist is in NP.*

Proof. The proof is based on a classical NP algorithm:

1. non-deterministically guess a set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$,
2. check in polynomial time (Proposition 10) whether it is a (weak or strong) stable extension.

Hence the result. □

4.1.4 Upper Bounds for Acceptability

First, consider the case of cf_x, for $x \in \{w, s\}$. An argument a is credulously accepted w.r.t. cf_x iff $\{a\} \in \mathsf{cf}_x(\mathcal{I})$. This can be easily checked, by verifying that $(a, a) \notin \mathcal{R}$ and $(a, a) \notin \mathcal{R}^?$. This is doable in polynomial time and logarithmic space. Thus cf_x-$\mathsf{Cred} \in \mathsf{L}$, for $x \in \{w, s\}$. Skeptical acceptability is even easier: since $\emptyset$ is weakly (respectively strongly) conflict-free, there is no skeptically acceptable argument w.r.t. cf_x for any IAF. The reasoning is the same for ad_x-Skep.

Proposition 13. *For $\sigma \in \{\mathsf{ad}, \mathsf{co}, \mathsf{st}, \mathsf{pr}\}$ and $x \in \{w, s\}$, σ_x-Cred is in NP.*

Proof. For $\sigma \in \{\mathsf{ad}, \mathsf{co}, \mathsf{st}\}$, guess a set of arguments that contains the queried argument a, and check (in polynomial time, see Proposition 10) whether it is a x-σ-extension. This is a NP algorithm for deciding σ_x-Cred.

For $\sigma = \mathsf{pr}$, notice that an argument belongs to some weakly (respectively strongly) preferred extension iff it belongs to some weakly (respectively strongly) admissible set, hence the result. $\square$

Proposition 14. *For* $\sigma \in \{\mathsf{co}, \mathsf{st}\}$ *and* $x \in \{w, s\}$*,* σ_x-Skep *is in* coNP*.*

Proof. Guess a set of arguments that does not contain the queried argument a and check (in polynomial time) whether it is a x-σ-extension, *i.e.* a is not skeptically accepted w.r.t. σ_x. This is a NP algorithm, thus σ_x-Skep is in coNP. $\square$

Proposition 15. *For* $x \in \{w, s\}$*,* pr_x-Skep *is in* Π_2^{P}*.*

Proof. Analogous to the proof of Proposition 14, except that the higher complexity of verification under the (weakly or strongly) preferred semantics (recall Proposition 11) yields a higher complexity upper bound for skeptical acceptability as well. $\square$

4.1.5 Upper Bounds for Non-Emptiness

Finally, we focus on the non-emptiness problem. We prove that it is doable in polynomial time and logarithmic space for (weak and strong) conflict-freeness, and NP-complete for other semantics. Again, NP-hardness results are implied by Proposition 9.

Proposition 16. cf_x-NE *is in* L*, for* $x \in \{w, s\}$*.*

Proof. First consider weak conflict-freeness. If a set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is weakly conflict-free, then every singleton $\{a\} \subseteq S$ is weakly conflict-free as well, so we focus on singletons. If $a \in \mathcal{A}^?$, then $\{a\}$ is trivially weakly conflict-free, so if $\mathcal{A}^? \neq \emptyset$, the answer is "YES". This can be checked in polynomial time and logarithmic space. Now assume $\mathcal{A}^? = \emptyset$, *i.e.* we consider singletons $\{a\}$ with $a \in \mathcal{A}$. $\{a\}$ is weakly conflict-free iff $(a, a) \notin \mathcal{R}$, which can be checked in polynomial time and logarithmic space.

Now, considering strong conflict-freeness, we can also focus on singletons, and simply check (again, in polynomial time and logarithmic space) that at least one argument $a \in \mathcal{A} \cup \mathcal{A}^?$ is not (certainly or uncertainly) self-attacking, *i.e.* $(a, a) \notin \mathcal{R} \cup \mathcal{R}^?$. $\square$

Proposition 17. σ_x-NE *is in* NP*, for* $\sigma \in \{\mathsf{ad}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$ *and* $x \in \{w, s\}$*.*

Proof. First, we consider $\sigma \in \{\mathsf{ad}, \mathsf{st}, \mathsf{co}\}$, and solve the problem with a simple NP algorithm:

1. non-deterministically guess a non-empty set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$.
2. check (in polynomial time, see Proposition 10) whether $S \in \sigma_x(\mathcal{F})$.

Hence the result for these semantics. Then, observe that the existence of a non-empty (weak or strong) admissible set implies the existence of a non-empty (weak or strong) preferred extension, which means that pr_x-NE has the same complexity as ad_x-NE. □

4.1.6 Discussion

Our complexity results are summarized in Table 5. We have proved that, in spite of the higher expressivity of IAFs compared to standard AFs, the complexity of most classical reasoning tasks remains the same. The only exception is skeptical acceptability under (weakly or strongly) complete semantics, for which we only have a coNP upper bound, while it is polynomial in standard Dung's AFs. We plan to study a counterpart of the grounded semantics for IAFs, which could bring new insights for the complete semantics. Finally, notice that using the weak or strong counterpart of our semantics does not have an impact on the complexity of reasoning.

Semantics σ_x	σ_x-Ver	σ_x-Cred	σ_x-Skep	σ_x-Exist	σ_x-NE
cf_x	in L	in L	trivial	trivial	in L
ad_x	in L	NP-c	trivial	trivial	NP-c
st_x	in L	NP-c	coNP-c	NP-c	NP-c
co_x	in L	NP-c	in coNP	trivial	NP-c
pr_x	coNP-c	NP-c	Π_2^P-c	trivial	NP-c

Table 5: Complexity of σ_x-Ver, σ_x-Exist, σ_x-Cred, σ_x-Skep and σ_x-NE for $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$ and $x \in \{w, s\}$. $\mathcal{C}$-c means $\mathcal{C}$-complete.

4.2 SAT-based Computational Approach

We follow the classical approach, initiated by [11], which consists in associating an AF with a propositional formula such that there is a bijection between the extensions of the AF and the models of the formula. Its has been applied with success for developing argumentation solvers [32, 41].

In the following, we consider an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, and we define a set of propositional variables $X_{\mathcal{A} \cup \mathcal{A}^?} = \{x_a \mid a \in \mathcal{A} \cup \mathcal{A}^?\}$. Intuitively, an interpretation ω corresponds to the set of arguments $S = \{a \in \mathcal{A} \cup \mathcal{A}^? \mid \omega(x_a) = \top\}$. We will provide in the rest of this section propositional formulas such that their models correspond to desirable sets of arguments (*e.g.* weakly or strongly conflict-free sets or extensions). This means that obtaining one (or each) extension can be done thanks to a SAT solver, providing one (or each) model of the formula. Credulous acceptance of an argument a can be checked by verifying that $\phi \wedge x_a$ is satisfiable (where ϕ is the formula corresponding to the chosen semantics), and skeptical acceptance corresponds to the unsatisfiability of $\phi \wedge \neg x_a$.

Conflict-freeness Recall that a set of arguments is weakly conflict-free if there is no certain attack between two certain arguments in it, while it is strongly conflict-free if there is no attack at all (neither certain nor uncertain) between any element of the set. This is encoded, respectively, by the following formulas ϕ^w_{cf} and ϕ^s_{cf}:

$$\phi^w_{\mathsf{cf}} = \bigwedge_{a,b \in \mathcal{A}, (a,b) \in \mathcal{R}} (\neg x_a \vee \neg x_b)$$

$$\phi^s_{\mathsf{cf}} = \bigwedge_{a,b \in \mathcal{A} \cup \mathcal{A}^?, (a,b) \in \mathcal{R} \cup \mathcal{R}^?} (\neg x_a \vee \neg x_b)$$

Admissibility Weak (respectively strong) admissibility is based on weak (respectively strong) conflict-freeness, and weak (respectively strong) defense. We introduce a formula δ_w (respectively δ_s) which characterizes sets of arguments that weakly (respectively strongly) defend all their elements.

$$\delta_w = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \to \bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

$$\delta_s = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \to \bigwedge_{b \in \mathcal{A} \cup \mathcal{A}^?, (b,a) \in \mathcal{R} \cup \mathcal{R}^?} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

Then, weak and strong admissibility are encoded in

$$\phi^x_{\mathsf{ad}} = \phi^x_{\mathsf{cf}} \wedge \delta_x$$

where $x \in \{w, s\}$.

Notice that δ_w and δ_s are not directly written as CNF formulas, but can be easily translated into ones:

$$x_a \to \bigwedge_{b \in X, (b,a) \in Y} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c \equiv \bigwedge_{b \in X, (b,a) \in Y} \neg x_a \vee \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

where X and Y are $\mathcal{A}$ and $\mathcal{R}$ (for δ_w) or $\mathcal{A} \cup \mathcal{A}^?$ and $\mathcal{R} \cup \mathcal{R}^?$ (for δ_s).

Complete Extensions The formulas δ_w and δ_s characterize sets of arguments that (weakly or strongly) defend all their elements. To characterize complete extensions, we just need to replace the implication by an equivalence, which yields sets of arguments that defend all their elements and contain everything they defend. Formally,

$$\phi^x_{\mathsf{co}} = \phi^x_{\mathsf{cf}} \wedge \delta'_x$$

where $x \in \{w, s\}$, and

$$\delta'_w = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \leftrightarrow \bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

$$\delta'_s = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \leftrightarrow \bigwedge_{b \in \mathcal{A} \cup \mathcal{A}^?, (b,a) \in \mathcal{R} \cup \mathcal{R}^?} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

For translating these formulas into CNF, we add auxiliary variables, using a technique similar to the one from [32]. For a given argument a, define y_a as a variable that is true when one of the certain attackers of a is accepted. This is encoded in the formulas by:

$$\bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} y_a \leftrightarrow \bigvee_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} x_b$$

which is easily translated into CNF, since

$$y_a \leftrightarrow \bigvee_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} x_b \equiv (\neg y_a \vee \bigvee_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} x_b) \wedge (\bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \neg x_b \vee y_a)$$

Then, δ'_w and δ'_s can be written in CNF:

$$x_a \leftrightarrow \bigwedge_{b \in X, (b,a) \in Y} y_b \equiv (\bigwedge_{b \in X, (b,a) \in Y} \neg x_a \vee y_b) \wedge (x_a \vee \bigvee_{b \in X, (b,a) \in Y} \neg y_b)$$

where, again, either $X = \mathcal{A}$ and $Y = \mathcal{R}$, or $X = \mathcal{A} \cup \mathcal{A}^?$ and $Y = \mathcal{R} \cup \mathcal{R}^?$.

Stable Extensions We recall that weakly (respectively strongly) stable extensions are weakly (respectively strongly) conflict-free sets that attack all the certain arguments (respectively all the arguments) that they do not contain. Said otherwise, it means that an argument which is not attacked by (a certain argument in) the extension belongs to the extension. It can be characterized as follows:

$$\phi^w_{\mathsf{st}} = \phi^w_{\mathsf{cf}} \wedge \bigwedge_{a \in \mathcal{A}} ((\bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \neg x_b) \rightarrow x_a)$$

$$\phi^s_{\mathsf{st}} = \phi^s_{\mathsf{cf}} \wedge \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} ((\bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \neg x_b) \rightarrow x_a)$$

Preferred Extensions Finally, weakly and strongly preferred semantics cannot (under the usual assumptions of complexity theory) be directly encoded as propositional formulas, since the complexity of reasoning with weak and strong preferred semantics is higher than the complexity of Boolean satisfiability (especially, skeptical acceptability is Π_2^{P}-complete). However, other techniques related to propositional logic have been used in the past for computing preferred extensions, *e.g.* quantified Boolean formulas [22], maximal satisfiable subsets [32] or CEGAR (CounterExample Guided Abstraction Refinement) [41]. These techniques could be adapted for solving classical reasoning problems under weakly or strongly preferred semantics.

5 Implementation and Experimentation

5.1 Implementation Details

We have implemented the approach described in Section 4.2, and conducted an empirical evaluation to assess the scalability of the approach. More precisely, we have solved the problem of producing one extension of an IAF (called SE-σ_x, see *e.g.* [33, 34]) for the semantics σ_x with $\sigma \in \{\mathsf{st}, \mathsf{co}\}$ and $x \in \{w, s\}$. This corresponds to obtaining one model (with a SAT solver) of the formula ϕ_σ^x. Recall that other classical reasoning tasks can be performed with a SAT solver as well:

- EE-σ_x: enumerate all the σ_x extensions is done by enumerating the models of ϕ_σ^x;
- CE-σ_x: counting the σ_x extensions is done by counting the models of ϕ_σ^x;
- DC-σ_x: deciding whether a given argument a is credulously accepted with respect to the semantics σ_x is done by checking whether $\phi_\sigma^x \wedge x_a$ is satisfiable;
- DS-σ_x: deciding whether a given argument a is skeptically accepted with respect to the semantics σ_x is done by checking whether $\phi_\sigma^x \wedge \neg x_a$ is unsatisfiable.

We can expect that DC-σ_x and DS-σ_x will be more or less equivalent to SE-σ_x regarding the hardness of practical resolution. Indeed, all these problems are solved by a single call to a SAT solver, with almost the same CNF as input of the SAT solver. On the contrary, enumeration and counting problems are in general much harder. These general intuitions about the relative difficulty of reasoning with abstract argumentation are in line with the results of the International Competition on Computational Models of Argumentation (ICCMA): we can observe that the scores of the solvers, during the last edition of the competition, are generally higher for the problems SE, DS and DC than for CE [34]. The same remark applies for the problem EE, as can be seen from the results of the 2017 edition of ICCMA [27].

We have implemented a Python script that reads a text file describing an IAF (using the same format as [43, 7]),[7] and produces the CNF encoding corresponding to ϕ_σ^x. This SAT instance is solved by the Python API PySAT [31], and the model obtained is then decoded in order to provide the extension to the user. Our code and its documentation are available online.[8]

5.2 Benchmark Generation and Experimental Setup

The goal of this preliminary experimentation is to assess the scalability of the approach, and to observe whether some parameters may influence the runtime.

[7] See `https://bitbucket.org/andreasniskanen/taeydennae/src/master/`.
[8] See `https://github.com/jgmailly/SAT-IAFs/`.

We have generated IAFs with the following method, based on Erdös-Rényi [23] (ER) graphs, *i.e.* random graphs built with two parameters: the number of nodes (n), and the probability, for two given nodes a and b, that there is an edge from a to b (p). This type of graph has been widely used in the literature on argumentation, including some past editions of the International Competition on Computational Models of Argumentation (ICCMA) [27]. For our experiments, we have used $n \in \{50, 100, 150, 200\}$ and $p \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$:

- For each (p, n), generate 5 AFs
- For each AF $\mathcal{F} = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$, generate four IAFs by selecting some arguments or attacks to be uncertain:
  - $\mathcal{I}_{1,1} = \langle \mathcal{A}^*, \emptyset, \mathcal{R}^*, \emptyset \rangle$, *i.e.* all arguments and attacks are certain;
  - $\mathcal{I}_{1,0.5} = \langle \mathcal{A}^*, \emptyset, \mathcal{R}, \mathcal{R}^? \rangle$ s.t. $\mathcal{R} \cup \mathcal{R}^? = \mathcal{R}^*$, and each attack in $\mathcal{R}^*$ has a probability 0.5 to be added to $\mathcal{R}^?$;
  - $\mathcal{I}_{0.5,1} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}^*, \emptyset \rangle$, s.t. $\mathcal{A} \cup \mathcal{A}^? = \mathcal{A}^*$, and each argument in $\mathcal{A}^*$ has a probability 0.5 to be added to $\mathcal{A}^?$;
  - $\mathcal{I}_{0.5,0.5} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ s.t. $\mathcal{R} \cup \mathcal{R}^? = \mathcal{R}^*$, and each attack in $\mathcal{R}^*$ has a probability 0.5 to be added to $\mathcal{R}^?$ and $\mathcal{A} \cup \mathcal{A}^? = \mathcal{A}^*$, and each argument in $\mathcal{A}^*$ has a probability 0.5 to be added to $\mathcal{A}^?$.

This generation model allows to test our SAT-based approach with various types of IAFs, and to observe whether the presence or absence of uncertain elements has an impact on runtime.

We have run the experiments on macOS 11.5, with a M1 Soc (3.2GHz) and 8GB of RAM.

5.3 Results

Tables 6 and 7 describe the results of the experiments for the (weak and strong) stable semantics and the (weak and strong) complete semantics, respectively. Lines correspond to pairs of values (p, n) used for generating the graphs, columns labeled by $\mathcal{I}_{X,Y}$ (with $X, Y \in \{1, 0.5\}$) correspond to the groups of instances defined in the previous section (with different ratios of arguments or attacks being uncertain), and the columns labeled "All" correspond to the union of the four previous groups. Reported numbers are the average runtime, rounded to 1 millisecond, for solving all the instances of one group.

The general observation made from this data is that our approach scales up well. Indeed, it solves the SE-σ_x problem in a few seconds in the worst case, for IAFs with 200 arguments (and even less than 1 second in most of cases).

Then, for each semantics in this experiment, we observe a correlation between the density of the graph (*i.e.* the probability p) and the runtime: a higher probability seems to imply a higher runtime (*ceteris paribus*). On the contrary, increasing the uncertainty seems to decrease the runtime: instances where half the arguments and half the attacks are uncertain ($\mathcal{I}_{0.5,0.5}$ groups) are

(p,n)	st_w $\mathcal{I}_{1,1}$	$\mathcal{I}_{1,0.5}$	$\mathcal{I}_{0.5,1}$	$\mathcal{I}_{0.5,0.5}$	All	st_s $\mathcal{I}_{1,1}$	$\mathcal{I}_{1,0.5}$	$\mathcal{I}_{0.5,1}$	$\mathcal{I}_{0.5,0.5}$	All
$(0.1,50)$	50	50	50	50	50	50	50	50	50	50
$(0.1,100)$	54	50	50	50	51	52	50	50	50	51
$(0.1,150)$	114	60	60	50	71	110	60	60	60	73
$(0.1,200)$	2388	296	70	60	704	2376	80	86	70	653
$(0.3,50)$	50	50	50	50	50	50	50	50	50	50
$(0.3,100)$	74	60	50	50	59	72	60	60	60	63
$(0.3,150)$	144	196	78	60	120	138	90	108	80	104
$(0.3,200)$	402	1920	106	84	628	382	154	186	124	212
$(0.5,50)$	50	50	50	50	50	50	50	50	50	50
$(0.5,100)$	80	70	60	50	65	80	70	74	70	74
$(0.5,150)$	162	150	90	70	118	152	120	132	106	128
$(0.5,200)$	320	676	146	102	311	294	214	258	180	237
$(0.7,50)$	50	50	50	50	50	50	50	50	50	50
$(0.7,100)$	100	76	66	60	76	90	80	88	76	84
$(0.7,150)$	200	172	106	80	140	190	140	164	130	156
$(0.7,200)$	404	440	186	126	289	370	270	320	230	298
$(0.9,50)$	60	50	50	50	53	50	50	50	50	50
$(0.9,100)$	110	82	70	60	81	110	90	100	80	95
$(0.9,150)$	246	164	122	96	157	226	170	200	152	187
$(0.9,200)$	510	320	226	146	300	464	330	400	290	371

Table 6: Average Runtime in Milliseconds for SE-st_x, $x \in \{w, s\}$

solved faster than instances with only arguments or only attack can be uncertain ($\mathcal{I}_{1,0.5}$ and $\mathcal{I}_{0.5,1}$), which in turn are solved faster than instance where all elements are certain ($\mathcal{I}_{1,1}$). Naturally, these conclusions are only preliminary, given the size of the benchmark. Further study will be conducted for determining whether the claim hold in general. It is interesting to notice that both cases where the average is higher than 1 second are very similar: they are the category of IAFs $\mathcal{I}_{1,1}$ with $p = 0.1$ and $n = 200$, respectively for st_w and st_s. For all the IAFs in this category, the runtime is higher than 1 second, and up to 3.5 seconds, which is (relatively) much higher than all the other runtimes. In particular, even for this category of graphs, runtimes for the variants of the complete semantics are below 0.2 second for each instance. Determining the cause of this difference is an interesting question for future work. Finally, let us conclude this analysis of the preliminary experiments by noticing that the chosen semantics (among $\{\mathsf{st}_w, \mathsf{st}_s, \mathsf{co}_w, \mathsf{co}_s\}$) has no strong impact (for the chosen benchmarks) on the average runtime.

6 Related Work

As mentioned in the introduction, most of the work on incomplete argumentation frameworks are strongly different, by nature, with our contribution, since

(p,n)	co_w $\mathcal{I}_{1,1}$	$\mathcal{I}_{1,0.5}$	$\mathcal{I}_{0.5,1}$	$\mathcal{I}_{0.5,0.5}$	All	co_s $\mathcal{I}_{1,1}$	$\mathcal{I}_{1,0.5}$	$\mathcal{I}_{0.5,1}$	$\mathcal{I}_{0.5,0.5}$	All
$(0.1,50)$	50	50	50	50	50	50	50	50	50	50
$(0.1,100)$	62	50	50	60	56	60	60	58	56	59
$(0.1,150)$	92	70	70	60	73	90	80	78	70	80
$(0.1,200)$	156	102	100	78	109	150	130	116	110	127
$(0.3,50)$	50	50	50	50	50	50	50	50	50	50
$(0.3,100)$	90	70	70	72	76	90	80	78	70	80
$(0.3,150)$	190	122	120	80	128	180	154	138	130	151
$(0.3,200)$	368	212	200	128	227	350	290	244	230	278
$(0.5,50)$	60	50	50	50	53	60	52	50	50	53
$(0.5,100)$	120	90	82	74	92	120	106	92	90	102
$(0.5,150)$	280	170	156	106	178	270	224	190	182	217
$(0.5,200)$	584	322	306	178	348	554	456	384	352	437
$(0.7,50)$	62	56	52	54	56	60	60	60	60	60
$(0.7,100)$	156	100	96	82	109	150	138	122	110	130
$(0.7,150)$	380	216	202	130	232	370	306	258	244	295
$(0.7,200)$	824	448	408	232	478	760	620	522	474	594
$(0.9,50)$	70	60	58	50	60	70	60	60	60	63
$(0.9,100)$	190	120	112	98	130	180	150	132	130	148
$(0.9,150)$	480	276	240	162	290	450	372	312	294	357
$(0.9,200)$	1028	558	530	302	605	960	780	658	602	750

Table 7: Average Runtime in Milliseconds for SE-co_x, $x \in \{w, s\}$

they rely on the set of completions of an IAF to define various decision problems [9, 7, 25, 26, 29, 38, 40]. Control Argumentation Frameworks (CAFs) [17, 37, 42] are highly related to IAFs. They add another kind of uncertainty (about the direction of an attack), and a "control part", which is a set of arguments and attacks that must be selected by the agent, the goal being to enforce the acceptability of a set of arguments in each (or some) completion, by means of the selected control arguments. Reasoning with CAFs is (again) only based on completions, and generally the computational complexity is high (at least the same as reasoning with completions of IAFs, and sometimes higher).

Reasoning with weighted AFs (*i.e.* AFs with weights on the attacks) [19] consists, somehow, in relaxing conflict-freeness in order to jointly accept conflicting arguments, as soon as the total amount of conflict (*i.e.* the sum of the weights of the attacks) is lower than a given inconsistency budget. We could adapt this principle for IAFs, by accepting only a given amount of conflict in extensions. Notice that weighted AFs, as defined in the literature, do not allow to distinguish between two types of arguments or attacks, which would be necessary to capture uncertain arguments or attacks. Other frameworks have "conflict tolerant" semantics, where sets of accepted arguments may not be conflict-free with respect to the initial attack relation, like Preference-based AFs [2], Valued-based AFs [10] or Strength-based AFs [44]. In all cases, the presence of conflicting arguments in the same extension can be explained by their relative priority: when the target has a higher priority than its attacker, then they can appear in the same extension. This is not the same intuition as ours, since only uncertainty justifies conflict-tolerance of our weak semantics.

The work by [28] shares some of the basic intuitions of our own contribution. The notions of conflict-freeness and self-defense are re-defined to take into account the number of attackers and defenders of arguments. Then, combining different types of conflict-freeness and self-defense induces different types of admissibility-based semantics. Moreover, this work introduces constraints to ensure that the semantics behave well (namely, that some extensions exist). These constraints can be reminiscent of the fact that, here, the notion of defense should be as strong as the notion of conflict-freeness to induce a "fundamental" (in the sense of Dung's fundamental lemma) notion of admissibility. However, the contribution of [28] focuses on standard AFs, *i.e.* no concept of uncertainty is involved. Adapting this approach to IAFs would then be an interesting research topic, that may provide intermediate solutions between our weak and strong semantics.

Finally, let us mention the work by [36] on Probabilistic Argumentation Frameworks (PrAFs). The relation between PrAFs and IAFs has already been discussed in the literature. Indeed, from the probability attached to arguments and attacks in a PrAF, one can deduce the probability of its completions (called induced AFs there), and then the probability that a set of arguments is an extension. So, reasoning with such PrAFs can be seen as a probabilistic extension of completion-based reasoning with IAFs [7, Section 8]. Intuitively, our direct approach for reasoning with IAFs (without relying on completions) could be extended to PrAFs as well, for instance, weak conflict-freeness could be parameterized by the probability of conflicts that can be tolerated in the set of arguments.

7 Conclusion

In this paper, we have continued an effort started by [14], and defined extension-based semantics for Incomplete Argumentation Frameworks (IAFs) that do not rely on the completions of the IAF. We have studied the properties of our new semantics, and provided complexity results and logical encodings for them. We have proven that the complexity of reasoning with these semantics is not harder than reasoning with classical extension-based semantics for abstract argumentation frameworks (in spite of the higher complexity of IAFs), and an experimental study shows that reasoning can be done efficiently thanks to modern SAT solving techniques.

Future work include, naturally, missing complexity results (*i.e.* tight results for the skeptical acceptability under weakly and strongly complete semantics) and a deeper experimental evaluation of the approach (*e.g.* using other types of graphs like Barabási-Albert [1], Watts-Strogatz [47] or thoses from the last ICCMA competition [34], or solving other problems than SE-σ). In particular, deeper experiments will allow to determine more accurately the impact of the number of uncertain elements on the runtime, as well as the impact of the choice of the semantics. Regarding the implementation, we have focused on the semantics that can be solved by a "simple" use of a SAT solver, *i.e.* such that the corresponding decision problem is at the first level of the polynomial hierarchy. The comparison of the various methods that reach the second level of the polynomial hierarchy (*e.g.* CEGAR-style algorithms [21], QBFs [22] or maximal satisfiable subsets [32]) for computing the preferred extensions is an enthralling question. The study of the grounded semantics will fulfill our study of Dung-style semantics for IAFs. Further theoretical results can be interesting, like *e.g.* a principle-based study in the spirit of [46]. A fundamental question concerns the weak stable semantics. The fact that weak stable extensions are not weakly admissible is quite surprising. The principle-based study will allow to determine whether the weak stable semantics satisfies interesting properties all the same, or whether an alternative definition is desirable (in particular, an alternative definition that would imply weak admissibility). We also plan to apply this kind of semantics to Control Argumentation Frameworks [17, 37], which would be a possible method to decrease the complexity of controllability. This requires to take into account the additional type of information, namely the uncertainty about the direction of attacks. The link with weighted AFs, *i.e.* integrating an inconsistency budget in the weak variants of our semantics, is also a promising line for future research.

Acknowledgement

The author thanks the reviewers of the preliminary conference version as well as the reviewers of this extended version for their valuable comments which helped to improve the paper.

References

[1] Réka Albert and Albert-Lázló Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.

[2] Leila Amgoud and Claudette Cayrol. A reasoning model based on the production of acceptable arguments. *Ann. Math. Artif. Intell.*, 34(1-3):197–215, 2002.

[3] Leila Amgoud, Claudette Cayrol, Marie-Christine Lagasquie-Schiex, and P. Livet. On bipolarity in argumentation frameworks. *Int. J. Intell. Syst.*, 23(10):1062–1093, 2008.

[4] Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009.

[5] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. Abstract argumentation frameworks and their semantics. In Pietro Baroni, Dov Gab-

bay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, pages 159–236. College Publications, 2018.

[6] Pietro Baroni, Massimiliano Giacomin, Beishui Liao, and Leon van der Torre. Encompassing uncertainty in argumentation schemes. In *Proceedings of the Workshop on Frontiers and Connections between Argumentation Theory and Natural Language Processing, Forlì-Cesena, Italy, July 21-25, 2014*, 2014.

[7] Dorothea Baumeister, Matti Järvisalo, Daniel Neugebauer, Andreas Niskanen, and Jörg Rothe. Acceptance in incomplete argumentation frameworks. *Artif. Intell.*, 295:103470, 2021.

[8] Dorothea Baumeister, Daniel Neugebauer, and Jörg Rothe. Verification in attack-incomplete argumentation frameworks. In *4th International Conference on Algorithmic Decision Theory (ADT 2015)*, pages 341–358, 2015.

[9] Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack. Verification in incomplete argumentation frameworks. *Artif. Intell.*, 264:1–26, 2018.

[10] Trevor J. M. Bench-Capon. Value-based argumentation frameworks. In Salem Benferhat and Enrico Giunchiglia, editors, *9th International Workshop on Non-Monotonic Reasoning (NMR 2002), April 19-21, Toulouse, France, Proceedings*, pages 443–454, 2002.

[11] Philippe Besnard and Sylvie Doutre. Checking the acceptability of a set of arguments. In *10th International Workshop on Non-Monotonic Reasoning (NMR 2004)*, pages 59–64, 2004.

[12] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2009.

[13] Martin Caminada and Paul E. Dunne. Strong admissibility revisited: Theory and applications. *Argument Comput.*, 10(3):277–300, 2019.

[14] Claudette Cayrol, Caroline Devred, and Marie-Christine Lagasquie-Schiex. Handling ignorance in argumentation: Semantics of partial argumentation frameworks. In *9th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2007)*, pages 259–270, 2007.

[15] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages 151–158. ACM, 1971.

[16] Sylvie Coste-Marquis, Caroline Devred, Sébastien Konieczny, Marie-Christine Lagasquie-Schiex, and Pierre Marquis. On the merging of Dung's argumentation systems. *Artif. Intell.*, 171(10-15):730–753, 2007.

[17] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Control argumentation frameworks. In *32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*, pages 4678–4685, 2018.

[18] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.

[19] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael J. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.*, 175(2):457–486, 2011.

[20] Wolfgang Dvorák and Paul E. Dunne. Computational problems in formal argumentation and their complexity. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, pages 631–688. College Publications, 2018.

[21] Wolfgang Dvorák, Matti Järvisalo, Johannes Peter Wallner, and Stefan Woltran. Complexity-sensitive decision procedures for abstract argumentation. *Artif. Intell.*, 206:53–78, 2014.

[22] Uwe Egly and Stefan Woltran. Reasoning in argumentation frameworks using quantified boolean formulas. In *Computational Models of Argument: Proceedings of COMMA 2006, September 11-12, 2006, Liverpool, UK*, volume 144, pages 133–144, 2006.

[23] Paul Erdös and Alfréd Rényi. On random graphs. i. *Publicationes Mathematicae*, pages 290–297, 1959.

[24] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 1712–1718, 2020.

[25] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over argument-incomplete aafs in the presence of correlations. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 189–195. ijcai.org, 2021.

[26] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over attack-incomplete aafs in the presence of correlations. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, pages 301–311, 2021.

[27] Sarah Alice Gaggl, Thomas Linsbichler, Marco Maratea, and Stefan Woltran. Design and results of the second international competition on computational models of argumentation. *Artif. Intell.*, 279, 2020.

[28] Davide Grossi and Sanjay Modgil. On the graded acceptability of arguments in abstract and instantiated argumentation. *Artif. Intell.*, 275:138–173, 2019.

[29] Andreas Herzig and Antonio Yuste-Ginel. Abstract argumentation with qualitative uncertainty: An analysis in dynamic logic. In Pietro Baroni, Christoph Benzmüller, and Yì N. Wáng, editors, *Logic and Argumentation - 4th International Conference, CLAR 2021, Hangzhou, China, October 20-22, 2021, Proceedings*, volume 13040 of *Lecture Notes in Computer Science*, pages 190–208. Springer, 2021.

[30] Anthony Hunter. Some foundations for probabilistic abstract argumentation. In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012*, volume 245 of *Frontiers in Artificial Intelligence and Applications*, pages 117–128. IOS Press, 2012.

[31] Alexey Ignatiev, Antonio Morgado, and Joao Marques-Silva. PySAT: A Python toolkit for prototyping with SAT oracles. In *SAT*, pages 428–437, 2018.

[32] Jean-Marie Lagniez, Emmanuel Lonca, and Jean-Guy Mailly. CoQuiAAS: A constraint-based quick abstract argumentation solver. In *27th IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2015, Vietri sul Mare, Italy, November 9-11, 2015*, pages 928–935. IEEE Computer Society, 2015.

[33] Jean-Marie Lagniez, Emmanuel Lonca, Jean-Guy Mailly, and Julien Rossit. Introducing the fourth international competition on computational models of argumentation. In *Proceedings of the Third International Workshop on Systems and Algorithms for Formal Argumentation co-located with the 8th International Conference on Computational Models of Argument (COMMA 2020), September 8, 2020*, volume 2672 of *CEUR Workshop Proceedings*, pages 80–85. CEUR-WS.org, 2020.

[34] Jean-Marie Lagniez, Emmanuel Lonca, Jean-Guy Mailly, and Julien Rossit. Design and results of ICCMA 2021. *CoRR*, abs/2109.08884, 2021.

[35] John Lawrence and Chris Reed. Argument mining: A survey. *Comput. Linguistics*, 45(4):765–818, 2019.

[36] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proceedings of the First International Workshop on Theory and Applications of Formal Argumentation (TAFA'11)*, pages 1–16, 2011.

[37] Jean-Guy Mailly. Possible controllability of control argumentation frameworks. In *8th International Conference on Computational Models of Argument (COMMA 2020)*, pages 283–294, 2020.

[38] Jean-Guy Mailly. Constrained incomplete argumentation frameworks. In Jirina Vejnarová and Nic Wilson, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 16th European Conference, ECSQARU 2021, Prague, Czech Republic, September 21-24, 2021, Proceedings*, volume 12897 of *Lecture Notes in Computer Science*, pages 103–116. Springer, 2021.

[39] Jean-Guy Mailly. Extension-based semantics for incomplete argumentation frameworks. In *Logic and Argumentation - 4th International Conference, CLAR 2021, Hangzhou, China, October 20-22, 2021, Proceedings*, pages 322–341, 2021.

[40] Jean-Guy Mailly. Yes, no, maybe, I don't know: Complexity and application of abstract argumentation with incomplete knowledge. *Argument Comput.*, 2022. To appear (doi: 10.3233/AAC-210010).

[41] Andreas Niskanen and Matti Järvisalo. μ-toksia: An efficient abstract argumentation reasoner. In Diego Calvanese, Esra Erdem, and Michael Thielscher, editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*, pages 800–804, 2020.

[42] Andreas Niskanen, Daniel Neugebauer, and Matti Järvisalo. Controllability of control argumentation frameworks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1855–1861, 2020.

[43] Andreas Niskanen, Daniel Neugebauer, Matti Järvisalo, and Jörg Rothe. Deciding acceptance in incomplete argumentation frameworks. In *34th AAAI Conference on Artificial Intelligence (AAAI 2020)*, pages 2942–2949, 2020.

[44] Julien Rossit, Jean-Guy Mailly, Yannis Dimopoulos, and Pavlos Moraitis. United we stand: Accruals in strength-based argumentation. *Argument Comput.*, 12(1):87–113, 2021.

[45] Gregory S. Tseytin. On the complexity of derivation in propositional calculus. *Studies in Constructive Mathematics and Mathematical Logic, Part II, Seminars in Mathematics*, pages 115–125, 1970.

[46] Leon van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation semantics. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leon van der Torre, editors, *Handbook of Formal Argumentation*, pages 797–837. College Publications, 2018.

[47] Duncan J. Watts and Steven Strogatz. Collective dynamics of "small-world" networks. *Nature*, 393:440–442, 1998.