

Yes, no, maybe, I don't know: Complexity and application of abstract argumentation with incomplete knowledge

Jean-Guy Mailly

LIPADE, Université de Paris, Paris, France

E-mail: jean-guy.mailly@u-paris.fr

Abstract. Abstract argumentation, as originally defined by Dung, is a model that allows the description of certain information about arguments and relationships between them: in an abstract argumentation framework (AF), the agent knows for sure whether a given argument or attack exists. It means that the absence of an attack between two arguments can be interpreted as “we know that the first argument does not attack the second one”. But the question of uncertainty in abstract argumentation has received much attention in the last years. In this paper, we survey approaches that allow to express information like “There may (or may not) be an attack between these arguments”. We describe the main models that incorporate qualitative uncertainty (or ignorance) in abstract argumentation, as well as some applications of these models. We also highlight some open questions that deserve some attention in the future.

Keywords: Abstract argumentation, uncertainty, ignorance, complexity

1. Introduction

Abstract argumentation [44] has become a prominent subfield of non-monotonic reasoning. The classical model consists of a directed graph where the nodes are pieces of information called *arguments*, and the edges represent the conflicts between arguments, called *attacks*. It can be applied to different domains, especially in multi-agent systems where argument-based dialogues can be used for deliberation [6], negotiation [4] or persuasion [27]. However, this model does not allow to express uncertain information, like “this argument may exist (or not)” or “there may be an attack from this argument to this one”. The increase of expressivity of the model (regarding uncertainty) has been motivated, first, by the issue of merging several argumentation frameworks (corresponding to the knowledge of different agents) into a global view of the group's knowledge [33]. In this context, Partial Argumentation Frameworks are defined as complete directed graphs with three different kinds of edges, that represent the certain existence of attacks, the certain non-existence of attacks, and the ignorance about the existence of attacks. The model expressivity has been further increased with the definition of Incomplete Argumentation Frameworks (IAFs) that allow to express uncertainty about the existence of arguments (with two different kinds of nodes in the graph). Finally, a new kind of uncertainty has been proposed in [39]: the ignorance about the direction of attacks. It has been defined in the context of Control Argumentation Frameworks, that allow to model uncertainty in dynamic and strategic scenarios, like automated negotiation [40,41].

The three kinds of uncertainty that have been considered in the literature can all be motivated by concrete situations. First, the uncertainty about the existence of an attack can be naturally identified with uncertainty about the preferences between arguments. Indeed, in Preference-based Argumentation Frameworks [2], the attack relation is combined with preferences over the set of arguments to build a new relation, generally called *defeat*. An argument a defeats an argument b if a attacks b and b is not preferred to a . Thus, if an agent is not sure whether b is preferred to a , she does not know whether a defeats b or not. An IAF can be used to abstract away from the Preference-based AF: pairs of arguments (a, b) such that a certainly defeat b can be considered as “certain attacks” in the IAF, while situations where the agent is not sure whether b is preferred to a are “uncertain attacks” in the IAF. For the same reason, if there are attacks (a, b) and (b, a) in the Preference-based AF, and the agent is not sure of the preferences between a and b , then it can be modeled as an “attack with uncertain direction” in the IAF: either a attacks b (modeling the case where a is preferred to b), or b attacks a (modeling the opposite situation), or both at the same time (modeling the situation where none of a and b is preferred to the other argument). Studying uncertainty about the existence of arguments has also good motivations. If we consider the underlying structure of arguments (*e.g.* logic-based arguments [24]), it makes sense to consider that, at some specific point in time, the premises of an argument may not be a consequence of an agent’s knowledge base (or the agent knows that these premises are revisable information). This means that this argument may not actually be “activated” for the moment (or will not remain so in the future). Also, even without considering the internal structure of arguments, uncertainty about the arguments is natural in scenarios like opponent modeling in dialogues. For instances, in argument-based negotiation [40,41], an agent may have a model of her opponent, that is used to help her choosing the best arguments. In such a situation, the agent does not usually know perfectly her opponent’s knowledge and beliefs, thus she cannot model with certainty which arguments the opponent will use in the dialogue.

This paper describes the main works in this line of research. After a brief description of background notions of propositional logic, computational complexity and abstract argumentation in Section 2, we describe Incomplete Argumentation Frameworks (IAFs, Section 3): the general model as well as subclasses and the complexity of reasoning for two kinds of approaches (based either on completions of the IAF, or on a re-definition of basic concepts like conflict-freeness or defense). Then, in Section 4, we present Control Argumentation Frameworks (CAFs), that allow the modeling of uncertainty in dynamic and strategic contexts. We discuss several applications of IAFs and CAFs in Section 5, and several related notions in Section 6. Finally, Section 7 concludes the paper and highlights some open questions.

2. Background

2.1. Propositional logic

We consider a set V of Boolean variables, *i.e.* variables that can be assigned a value in $\mathbb{B} = \{0, 1\}$, where 0 is interpreted as *false* and 1 as *true*. A well-formed propositional formula ϕ is then:

- $\phi = x$, with $x \in V$, is called an atomic formula,
- $\phi = \phi_1 \wedge \phi_2$, with ϕ_1 and ϕ_2 two well-formed formulas, is the conjunction of ϕ_1 and ϕ_2 ,
- $\phi = \phi_1 \vee \phi_2$, with ϕ_1 and ϕ_2 two well-formed formulas, is the disjunction of ϕ_1 and ϕ_2 ,
- $\phi = \neg\phi'$, with ϕ' a well-formed formula, is the negation of ϕ' .

We use $var(\phi)$ to denote the set of Boolean variables that appear in ϕ .

An interpretation is then a mapping $\omega : V \rightarrow \mathbb{B}$, *i.e.* an assignment of a truth value to every variable. The notion can be generalized to formulas:

- $\omega(\phi_1 \wedge \phi_2) = 1$ if and only if $\omega(\phi_1) = 1$ and $\omega(\phi_2) = 1$,
- $\omega(\phi_1 \vee \phi_2) = 1$ if and only if $\omega(\phi_1) = 1$ or $\omega(\phi_2) = 1$,
- $\omega(\neg\phi) = 1$ if and only if $\omega(\phi) = 0$.

An interpretation ω is called a *model* of ϕ if $\omega(\phi) = 1$. The set of models of ϕ is denoted by $\text{mod}(\phi)$. Determining whether a formula has at least one model is the well-known Boolean Satisfiability problem (SAT). Although it is theoretically intractable (see Section 2.2), modern SAT solvers can handle large instances of this problem [54]. Two formulas ϕ, ψ are *logically equivalent* when they have the same set of models, which is denoted by $\phi \equiv \psi$. We use this notion to introduce additional connectives, that allow a more compact writing of formulas:

- the implication $\phi \rightarrow \psi \equiv \neg\phi \vee \psi$,
- the equivalence $\phi \leftrightarrow \psi \equiv (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$.

2.2. Computational complexity

In this part, we briefly describe some background notions on computational complexity, that are important for understanding the relative hardness of the computational problems that will be presented. First, let us mention that we focus on *decision problems*, that are computational problems that must be answered by a binary answer “YES” or “NO”. We will describe *complexity classes*, that are simply sets of decision problems that share some properties regarding the time (or space) necessary for solving them. The first one is made of problems that are considered to be *tractable*: \mathbf{P} is the set of decision problems that can be solved in polynomial time. More formally, for any $\mathcal{P} \in \mathbf{P}$, there is a deterministic algorithm that can decide any instance i of \mathcal{P} in (at most) n^k computation steps, where n is the size of the instance i , and $k \in \mathbb{N}$ is a fixed constant. The assumption that the algorithm is deterministic is important here. A subclass of \mathbf{P} that we consider in this paper is \mathbf{L} , the set of decision problems that can be decided in *logarithmic space* by a deterministic algorithm.

While all these problems are considered as “easy” ones, we also define classes characterized by non-deterministic algorithms that are considered as computationally hard. Let us first introduce \mathbf{NP} , the set of decision problems that can be solved in polynomial time by a *non-deterministic algorithm*. Another way to characterize \mathbf{NP} is to use the notion of *witness*. Given an instance i of a problem $\mathcal{P} \in \mathbf{NP}$, define $w(i)$ the set of potential witnesses that prove that i is a “YES” instance of \mathcal{P} . \mathbf{NP} is the set of decision problems such that for each instance i ,

- $w(i)$ is non-empty,
- for any $x \in w(i)$, verifying whether x is actually a witness is in \mathbf{P} ,
- if i is actually a “YES” instance, then there is at least one $x \in w(i)$ that is an actual witness for i .

For clarifying the concept of witness, let us mention the classical Boolean Satisfiability problem (SAT). Potential witnesses that a propositional formula ϕ is satisfiable are all the propositional interpretations on the set of variables $\text{var}(\phi)$, so $w(\phi) \neq \emptyset$. For any propositional interpretation, it can be polynomially checked whether it satisfies ϕ or not, *i.e.* proving whether it is actually a witness for ϕ . A usual way to rephrase the definition of \mathbf{NP} using witnesses is the following one: \mathbf{NP} is the set of problems that may be hard to solve, but such that checking a solution is easy.

The complementary of a complexity class \mathcal{C} is $\overline{\mathcal{C}} = \{\overline{\mathcal{P}} \mid \mathcal{P} \in \mathcal{C}\}$, where the complementary problem $\overline{\mathcal{P}}$ of \mathcal{P} is the decision problem built on the same set of instances as \mathcal{P} , such that the “YES” instances become “NO” instances (and vice-versa). Then, coNP can be defined as $\overline{\text{NP}}$. It is well-known that $\text{P} \subseteq \text{NP}$ and $\text{P} \subseteq \text{coNP}$.¹

The classes P, NP and coNP form the first levels of the *polynomial hierarchy* [79]. To define higher levels, we need to introduce the concept of *oracle*. An oracle for a class \mathcal{C} can be seen as a black box that solves a problem from \mathcal{C} in a single step. Then, the complexity class $\mathcal{C}^{\mathcal{C}'}$ is the set of decision problems that can be solved by a \mathcal{C} algorithm that has access to a \mathcal{C}' oracle. For instance, P^{NP} is the set of problems that can be solved in polynomial time by a deterministic algorithm that has access to an NP oracle. The concept of oracles allows to define recursively the classes Δ_k^{P} , Σ_k^{P} and Π_k^{P} , where $k \in \mathbb{N}$. Formally,

$$\begin{cases} \Delta_0^{\text{P}} = \Sigma_0^{\text{P}} = \Pi_0^{\text{P}} = \text{P} \\ \forall k > 0, \Delta_k^{\text{P}} = \text{P}^{\Sigma_{k-1}^{\text{P}}} \\ \forall k > 0, \Sigma_k^{\text{P}} = \text{NP}^{\Sigma_{k-1}^{\text{P}}} \\ \forall k > 0, \Pi_k^{\text{P}} = \text{coNP}^{\Sigma_{k-1}^{\text{P}}} \end{cases}$$

These classes are related by inclusion relations: for any $k \in \mathbb{N}$, $\Delta_k^{\text{P}} \subseteq \Sigma_k^{\text{P}}$ and $\Delta_k^{\text{P}} \subseteq \Pi_k^{\text{P}}$. In turn, $\Sigma_k^{\text{P}} \subseteq \Delta_{k+1}^{\text{P}}$ and $\Pi_k^{\text{P}} \subseteq \Delta_{k+1}^{\text{P}}$. These inclusions are illustrated by Fig. 1, where an arrow from \mathcal{C} to \mathcal{C}' means that $\mathcal{C} \subseteq \mathcal{C}'$. It is usually assumed that the polynomial hierarchy does not collapse (*i.e.* the inclusions are strict), although it has never been proven. This means that problems in a complexity class from the level k are supposed to be harder than problems from a level $k' < k$.

We present a last complexity class, namely DP, which stands for *Differential Polynomial Time*. It was first introduced by [75], as the set of decision problems

$$\text{DP} = \{\mathcal{P}_1 \cap \mathcal{P}_2 \mid \mathcal{P}_1 \in \text{NP}, \mathcal{P}_2 \in \text{coNP}\}$$

where $\mathcal{P}' = \mathcal{P}_1 \cap \mathcal{P}_2$ is the decision problem built from the same instances as \mathcal{P}_1 and \mathcal{P}_2 , and $i \in \mathcal{P}_1 \cap \mathcal{P}_2$ is a “YES” instance of \mathcal{P}' if and only if it is a “YES” instance for both \mathcal{P}_1 and \mathcal{P}_2 . Both NP and coNP are included in DP, which is in turn included in Δ_2^{P} .

Let us conclude this part on computational complexity with the notion of *hardness*. Intuitively, a problem \mathcal{P} is said to be \mathcal{C} -hard if it is as hard as any problem $\mathcal{P}' \in \mathcal{C}$. The notion of “as hard as” is formally defined through the notion of *reduction*. A reduction from \mathcal{P}' to \mathcal{P} is a mapping $f : I_{\mathcal{P}'} \rightarrow I_{\mathcal{P}}$ (where $I_{\mathcal{P}'}$ and $I_{\mathcal{P}}$ are, respectively, the instances of \mathcal{P}' and the instances of \mathcal{P}), such that $i \in I_{\mathcal{P}'}$ is a




Fig. 1. The polynomial hierarchy.

¹Whether the opposite inclusions hold or not is still an open problem at the time of writing. See *e.g.* <https://www.claymath.org/millennium-problems/p-vs-np-problem>.

“YES” instance of \mathcal{P}' if and only if $f(i)$ is a “YES” instance of \mathcal{P} . If f can be computed in polynomial time, then we write $\mathcal{P}' \leq_f^P \mathcal{P}$, which means that \mathcal{P} is (at least) as hard as \mathcal{P}' . For \mathcal{C} a class of the polynomial hierarchy with $k \geq 1$, we say that \mathcal{P} is \mathcal{C} -hard if and only if $\forall \mathcal{P}' \in \mathcal{C}, \mathcal{P}' \leq_f^P \mathcal{P}$. A notion of P-hardness can also be defined, using logarithmic reductions instead of polynomial ones, *i.e.* f must be computable in polynomial time, using logarithmic space. While the notion of hardness provides a lower bound for the complexity of a problem, the membership of a problem \mathcal{P} to a class \mathcal{C} provides an upper bound. When both bounds coincide (*i.e.* \mathcal{P} is \mathcal{C} -hard, and $\mathcal{P} \in \mathcal{C}$), we say that \mathcal{P} is \mathcal{C} -complete. This means, intuitively, that \mathcal{P} is one of the hardest problems in \mathcal{C} . Since \leq_f^P (and the logarithmic counterpart) is transitive, there is no need to prove that $\mathcal{P}' \leq_f^P \mathcal{P}$ for each $\mathcal{P}' \in \mathcal{C}$ for showing that \mathcal{P} is \mathcal{C} -hard; proving that $\mathcal{P}' \leq_f^P \mathcal{P}$ for some \mathcal{C} -hard problem \mathcal{P}' is enough. SAT (and variants of it) can be used for that, since it is the first problem that was proven to be NP-complete [32].

We refer the interested reader to [7] for more details on the theory of computational complexity.

2.3. Abstract argumentation

Abstract argumentation has been introduced by Dung [44] as the study of arguments acceptability based only on their relationship. It means that arguments are considered as atomic pieces of information: their internal structure and their origin are ignored in the reasoning process. Only the attacks between them matter for determining which arguments can be accepted. Formally:

Definition 1 (Argumentation Framework [44]). An abstract Argumentation Framework (AF) is a directed graph $\mathcal{F} = \langle A, R \rangle$ where the nodes A are *arguments*, and the edges $R \subseteq A \times A$ are *attacks*.

In this paper, we consider only *finite* AFs, *i.e.* A is a finite set of arguments. For $a, b \in A$, we say that a *attacks* b if $(a, b) \in R$. The notion of attack can be extended to sets of arguments: $S \subseteq A$ attacks an argument $b \in A$ if there is some $a \in S$ that attacks b . Finally, an argument $a \in A$ is *defended* by the set $S \subseteq A$ if, $\forall b \in S$ such that $(b, a) \in R, \exists c \in S$ such that $(c, b) \in R$ (*i.e.* each attacker of a is counter-attacked by an element of S).

Example 1. Figure 2 depicts an AF $\mathcal{F} = \langle A, R \rangle$ with $A = \{a_1, a_2, a_3, a_4, a_5, a_6, a_7\}$ and $R = \{(a_1, a_2), (a_2, a_3), (a_3, a_4), (a_4, a_3), (a_4, a_5), (a_5, a_6), (a_6, a_7), (a_7, a_5)\}$. The set of arguments $S = \{a_1, a_3\}$ defends a_3 since both its attackers are attacked by an element of S : we have $(a_1, a_2) \in R$ and $(a_3, a_4) \in R$.

The acceptability of arguments is classically evaluated through the notion of *extensions*, *i.e.* sets of collectively acceptable arguments. A set of arguments has to satisfy some properties to be an extension. The basic principles that underly most extension-based semantics are conflict-freeness and admissibility:




Fig. 2. An example of argumentation framework.

Definition 2 (Conflict-freeness and Admissibility). Given $\mathcal{F} = \langle A, R \rangle$ an AF, $S \subseteq A$ is called *conflict-free* if $\forall a, b \in S, (a, b) \notin R$. A conflict-free set is called *admissible* if it defends all its elements.

Conflict-free and admissible sets of an AF \mathcal{F} are (respectively) denoted $cf(\mathcal{F})$ and $adm(\mathcal{F})$. Conflict-freeness means that a set of arguments must be internally coherent in order to be acceptable, and admissibility means that it must somehow survive any threat, *i.e.* any attack against it.

Example 2. Continuing the previous example, observe that the set of arguments $S = \{a_1, a_3\}$ is admissible: it is conflict-free, and it defends all its elements.

From these basic requirements, extensions are defined following four semantics, proposed in [44]. Formally, an extension-based semantics is a mapping from an AF $\mathcal{F} = \langle A, R \rangle$ to a set of sets of arguments $\sigma(\mathcal{F}) \subseteq 2^A$.

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

- a *complete extension* (*co*) if it contains all the arguments that it defends;
- a *preferred extension* (*pr*) if it is a \subseteq -maximal admissible set;
- a *grounded extension* (*gr*) if it is a \subseteq -minimal complete extension;
- a *stable extension* (*stb*) if it attacks all the arguments in $A \setminus S$.

Other semantics have been defined since then; see [9] for an overview. It is well-known since Dung's seminal paper [44] that, for any AF \mathcal{F} , $|gr(\mathcal{F})| = 1$, $stb(\mathcal{F}) \subseteq pr(\mathcal{F}) \subseteq co(\mathcal{F})$, and $|pr(\mathcal{F})| \geq 1$. On the contrary, the stable semantics is not universally defined, *i.e.* there are some AFs \mathcal{F} with $stb(\mathcal{F}) = \emptyset$.

From a set of extensions $\sigma(\mathcal{F})$, two classical reasoning modes are used to determine the set of acceptable arguments.

Definition 4 (Argument Acceptability). Given $\mathcal{F} = \langle A, R \rangle$ an AF and σ an extension-based semantics, an argument $a \in A$ is

- *credulously acceptable* with respect to σ if it belongs to some σ -extension of \mathcal{F} ,
- *skeptically acceptable* with respect to σ if it belongs to each σ -extension of \mathcal{F} ,
- *rejected* otherwise.

We use $cred_\sigma(\mathcal{F})$ and $skep_\sigma(\mathcal{F})$ to denote the sets of credulously and skeptically acceptable arguments. Under most semantics, $skep_\sigma(\mathcal{F}) \subseteq cred_\sigma(\mathcal{F})$ and the rejected arguments are $A \setminus cred_\sigma(\mathcal{F})$. Notice that it is not the case under the stable semantics, for AFs \mathcal{F} such that $stb(\mathcal{F}) = \emptyset$. In that case, $cred_{stb}(\mathcal{F}) = \emptyset$, and $skep_{stb}(\mathcal{F}) = A$.

Example 3. We continue the previous example again. Considering \mathcal{F} from Fig. 2, and $\sigma \in \{gr, stb, pr, co\}$, Table 1 describes the σ -extensions of \mathcal{F} , as well as the sets of credulously and skeptically acceptable arguments.

The complexity of various reasoning problems has been studied in the past. A recent overview is provided by [48]. Let us formally define some of these problems, namely *Verification* (Ver), *Credulous Acceptability* (CA), *Skeptical Acceptability* (SA) and *Non-Empty Existence* (NE), that will be of interest in the rest of the paper.

Table 1

Extensions of \mathcal{F} and acceptable arguments under $\sigma \in \{gr, stb, pr, co\}$

Semantics σ	$\sigma(\mathcal{F})$	$cred_\sigma(\mathcal{F})$	$skep_\sigma(\mathcal{F})$
<i>gr</i>	$\{\{a_1\}\}$	$\{a_1\}$	$\{a_1\}$
<i>stb</i>	$\{\{a_1, a_4, a_6\}\}$	$\{a_1, a_4, a_6\}$	$\{a_1, a_4, a_6\}$
<i>pr</i>	$\{\{a_1, a_4, a_6\}, \{a_1, a_3\}\}$	$\{a_1, a_3, a_4, a_6\}$	$\{a_1\}$
<i>co</i>	$\{\{a_1, a_4, a_6\}, \{a_1, a_3\}, \{a_1\}\}$	$\{a_1, a_3, a_4, a_6\}$	$\{a_1\}$

Table 2

Complexity of σ -Ver, σ -CA, σ -SA and σ -NE, for $\sigma \in \{cf, adm, gr, stb, co, pr\}$ [34,42,44,45,49]. C-c means C-complete

Semantics σ	σ -Ver	σ -CA	σ -SA	σ -NE
<i>cf</i>	in L	in L	trivial	in L
<i>adm</i>	in L	NP-c	trivial	NP-c
<i>gr</i>	P-c	P-c	P-c	in L
<i>stb</i>	in L	NP-c	coNP-c	NP-c
<i>co</i>	in L	NP-c	P-c	NP-c
<i>pr</i>	coNP-c	NP-c	Π_2^P -c	NP-c

 σ -Ver Given $\mathcal{F} = \langle A, R \rangle$ an AF and $S \subseteq A$, is S a σ -extension of \mathcal{F} ? σ -CA Given $\mathcal{F} = \langle A, R \rangle$ an AF and $a \in A$, is a credulously acceptable in \mathcal{F} under σ ? σ -SA Given $\mathcal{F} = \langle A, R \rangle$ an AF and $a \in A$, is a skeptically acceptable in \mathcal{F} under σ ? σ -NE Given $\mathcal{F} = \langle A, R \rangle$ an AF, does \mathcal{F} admit a non-empty σ -extension?

Their complexity is described in Table 2. In the table, “trivial” means that the answer is trivially “NO” for any instance: indeed, \emptyset is conflict-free and admissible for every AF, so there is no skeptically acceptable argument with respect to conflict-freeness and admissibility.

Finally, let us discuss the specificity of some semantics regarding the empty set and skeptical acceptability. The stable semantics is the only semantics for which there are AFs \mathcal{F} such that $stb(\mathcal{F}) = \emptyset$. We have already mentioned that, in this case, the classical definition of skeptical acceptability leads to the acceptability of each argument. This may seem counterintuitive for some applications. This has conducted to the definition of an alternative decision problem, which consists in determining whether the AF has at least one extension (here, the non-emptiness is not important), and then whether the given argument is skeptically accepted. This is the problem of *Existence and Skeptical Acceptability* (ExSA), formally:

 σ -ExSA Given $\mathcal{F} = \langle A, R \rangle$ an AF and $a \in A$, do $\sigma(\mathcal{F}) \neq \emptyset$ and $a \in \bigcap_{E \in \sigma(\mathcal{F})} E$ hold?

For the other semantics from Definition 3, (as well as conflict-freeness and admissibility), the existence of at least one extension for any AF is guaranteed, so ExSA coincides with SA. However, the empty set yields some problem for skeptical acceptability in other cases. Indeed, as mentioned before, σ -SA is trivial (with all instances being “NO” instances) for $\sigma \in \{cf, adm\}$, because \emptyset is conflict-free and admissible for any AF. This issue has conducted to the definition of alternative versions of these “semantics”: $\sigma_{\neq \emptyset}$, where $\sigma \in \{cf, adm\}$, is defined by $\sigma_{\neq \emptyset}(\mathcal{F}) = \sigma(\mathcal{F}) \setminus \{\emptyset\}$. These alternative versions of conflict-freeness and admissibility solve the issue of triviality for skeptical acceptance, but they face the same issue as the stable semantics, since $\sigma_{\neq \emptyset}(\mathcal{F})$ may be empty. Thus, ExSA has also been studied for these semantics. The complexity of ExSA is described in Table 3.

Table 3

Complexity of σ -ExSA for $\sigma \in \{cf_{\neq\emptyset}, adm_{\neq\emptyset}, stb\}$ [47,73]

Semantics σ	σ -ExSA
$cf_{\neq\emptyset}$	in P
$adm_{\neq\emptyset}$	DP-complete
stb	DP-complete

3. Incomplete argumentation frameworks

In this section, we describe the first models of incompleteness in abstract argumentation, and then Incomplete Argumentation Frameworks (IAFs) that generalize both previous models. We present the complexity and algorithms for several related decision problems. We also present a particular family of semantics, initially defined for Partial AFs, a subclass of IAFs, and then generalized to the full IAF model.

3.1. Attack-incompleteness and argument-incompleteness

From a historical point of view, the first model that extends Dung's abstract argumentation framework with incompleteness is the so-called Partial Argumentation Framework (PAF), proposed in [33]. A PAF is a tuple $\mathcal{P} = \langle A, R, I, N \rangle$ where A is a set of arguments, and R, I, N form a partition of $A \times A$, where R is the *attack* relation, I is the *ignorance* relation, and N is the *non-attack* relation. Intuitively, $(a, b) \in R$ means that the agent knows for sure that a attacks b , while $(a, b) \in I$ means that the agent ignores whether it is the case or not. Finally, $(a, b) \in N$ means that the agent knows for sure that a does not attack b . As mentioned by the authors of the original paper, N can be deduced from R and I (since $N = (A \times A) \setminus (R \cup I)$), thus any PAF can be fully specified by $\langle A, R, I \rangle$. This model is used in a context of merging several argumentation frameworks, each of them representing the knowledge of one agent [33] (see Section 5.1 for more details). The notion of completion of a PAF $\langle A, R, I \rangle$ is defined as a standard AF that "solves" the uncertainty in the PAF, *i.e.* $\langle A, R' \rangle$ such that $R \subseteq R' \subseteq R \cup I$. The model was further studied in [29], where semantics adapted to this setting are defined (see Section 3.4), without using the notion of completion.

While the literature on PAFs is quite sparse, the model has been later re-named as Attack-Incomplete Argumentation Framework (Att-IAF) [16]. An Att-IAF is then a tuple $\langle A, R, R^? \rangle$ (where $R^?$ corresponds to I in PAFs). The focus of [16] is the verification problem, *i.e.* given an Att-IAF \mathcal{I} , a set of arguments S and a semantics σ , is S a σ -extension of some (or each) completion of \mathcal{I} ?

Then, a similar question has been studied for Argument-Incomplete Argumentation Framework (Arg-IAF) [19], *i.e.* argumentation frameworks where the actual existence of some arguments is uncertain. An Arg-IAF is a tuple $\langle A, A^?, R \rangle$ such that $R \subseteq A \cup A^?$, and a completion is then a standard AF $\langle A', R' \rangle$ where $A \subseteq A' \subseteq A \cup A^?$, and $R' = R \cap (A' \times A')$.

The (generalized) Incomplete Argumentation Framework models situations where both arguments and attacks can be uncertain. We focus on this framework in the rest of the section, and give more details on the reasoning based on completions (similar to the approach studied in [16,19]), and the reasoning based on redefinition of basic notions like conflict-freeness and defense (in the spirit of [29]).

3.2. The generalized IAF model

Incomplete AFs were first defined in [18], which generalizes the work from [16,19].

Definition 5 (Incomplete Argumentation Framework). An *Incomplete Argumentation Framework* (IAF) is a tuple $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ where A and $A^?$ are disjoint sets of arguments, and $R, R^? \subseteq (A \cup A^?) \times (A \cup A^?)$ are disjoint sets of attacks.

Elements in A and R are called *certain* (or *definite*) arguments and attacks, while $A^?$ and $R^?$ are *uncertain* arguments and attacks. We have already briefly discussed, in the introduction, the intuitive meaning of the different kinds of arguments and attacks. Roughly speaking, they can be seen as an abstraction of information regarding the internal structure of arguments (e.g. the validity of the arguments premises may change through time), or additional knowledge (e.g. unknown preferences between arguments). They can also represent potential evolution of the AF (e.g. in the case of a debate, uncertain arguments and attacks can be seen as those that can be used – or not – by other agents). Notice a particular case: an uncertain argument $a \in A^?$ that certainly attacks (or is attacked by) another argument $b \in A \cup A^?$, i.e. $(a, b) \in R$ (or $(b, a) \in R$). This specific kind of attack is sometimes called *conditionally definite*, which means that this attacks certainly exists *if* a and b actually exist. A concrete case where this happens is a debate, where a is an argument that could be used by another agent. We can be sure that – if it appears in the debate – a attacks b (e.g. because a thorough analysis of their internal content shows some logical inconsistency between them), but there is no certainty that the opponent will use a .

Example 4. Figure 3 describes an IAF $\mathcal{I} = \langle A, A^?, R, R^? \rangle$. Certain arguments $A = \{a, b, c\}$ and attacks $R = \{(b, a), (e, b)\}$ are represented by plain nodes and arrows, while uncertain attacks $R^? = \{(c, a), (d, b)\}$ are dotted arrows, and uncertain arguments $A^? = \{d, e\}$ are dashed nodes.

The (most classical) way of reasoning with an IAF is based on the notion of completions, that are (standard) AFs corresponding to a possible resolution of the uncertainty encoded in the IAF. Before formally defining completions, we introduce a notation: given a set of attacks Att and a set of arguments Arg , $Att|_{Arg} \triangleq Att \cap (Arg \times Arg)$.

Definition 6 (Completion of an IAF). Given an IAF $\mathcal{I} = \langle A, A^?, R, R^? \rangle$, a completion is an AF $\mathcal{F} = \langle A', R' \rangle$ with

- $A \subseteq A' \subseteq A \cup A^?$,
- $R|_{A'} \subseteq R' \subseteq (R \cup R^?)|_{A'}$.

The set of completions of \mathcal{I} is $comp(\mathcal{I})$.

Example 5. Figure 4 gives the completions of \mathcal{I} (Fig. 3). We can observe that certain arguments appear in all completions, and certain attacks appear in all the completions that contain both incident arguments (e.g. (e, b) appears in all the completions where e appears).

This model is quite general, and includes previously defined models (i.e. those described in Section 3.1) as subclasses. First, a *Partial AF* (PAF) [33] $\mathcal{P} = \langle A, R, I \rangle$ can be seen as an IAF with $A^? = \emptyset$


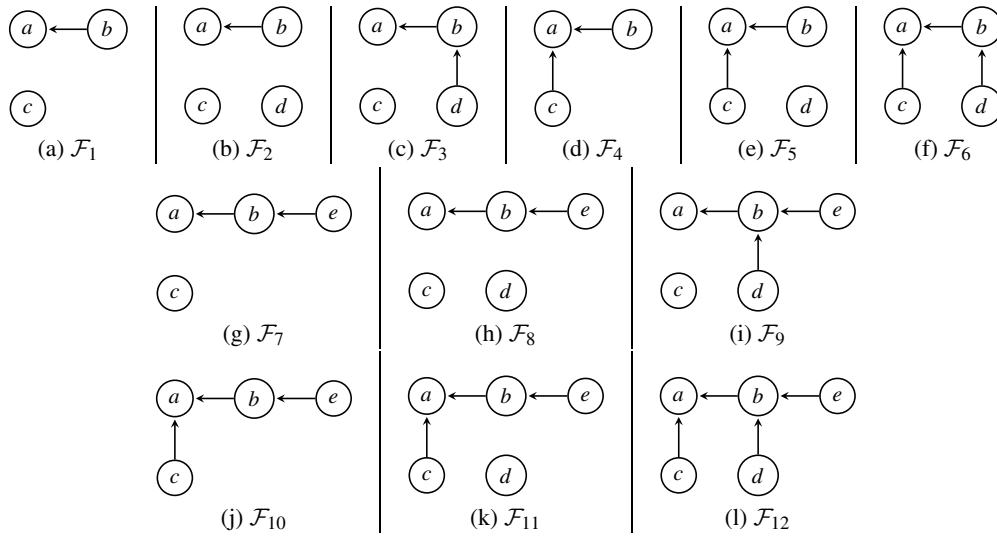


Fig. 3. An IAF \mathcal{I} .

Fig. 4. The completions of \mathcal{Z} .

and $R^? = I$. Recall that the same model was studied under the name *Attack Incomplete AF* (Att-IAF) in [16]. On the contrary, an *Argument Incomplete AF* (Arg-IAF) [19] is a tuple $\langle A, A^?, R \rangle$, *i.e.* an IAF with $R^? = \emptyset$.

Two families of reasoning approaches have been defined. In most of the cases, reasoning is based on the set of completions. For instance, an argument is possibly accepted if it is accepted in some completion, and necessarily accepted if it is accepted in each completion. We describe this approach in Section 3.3. Another approach (the first one, from a chronological point of view) consists in adapting the basic notions of conflict-freeness and defense, on which Dung's semantics are based. This allows to define new semantics suited to the specific framework at hand. It was proposed in [29] for PAFs and later generalized to IAFs [65], and we describe it in Section 3.4.

3.3. Possible and necessary reasoning

Similarly to credulous and skeptical reasoning, based on a set of extensions, two modes of reasoning can be used when inference is made from a set of completions. Possible reasoning is the counterpart (for completions) of credulous reasoning (for extensions). It means that we expect some property to be satisfied in *some* completion. On the contrary, necessary reasoning is the counterpart (for completions) of skeptical reasoning (for extension): the property must be satisfied in *each* completion. A natural example where both reasoning modes make sense is a trial. There may be uncertainty arising in such scenario. The prosecutor's goal is to prove that the defendant is guilty without any doubt, *i.e.* it must be a consequence of every completion. On the contrary, the defendant's lawyer has to prove that the defendant may be innocent (this is the presumption of innocence), *i.e.* it must be a consequence of some completion.

Possible and necessary modes of reasoning can be combined with any decision problem of interest regarding argumentation frameworks. Interestingly, being "skeptical" about the completions does not mean that the agent has to be skeptical about the extensions (and vice-versa). So, both credulous and skeptical reasoning, as well as verification and non-empty existence, have been adapted for IAFs under the possible and necessary views.

3.3.1. Complexity

Now we discuss the different reasoning problems for IAFs that have been studied in the literature, as well as their complexity.

Verification. The first study on extension verification for IAFs was described in [18]:²

σ -IncPV Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $S \subseteq A \cup A^?$, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ such that $S \cap A'$ is a σ -extension of \mathcal{F} ?

σ -IncNV Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $S \subseteq A \cup A^?$, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, is $S \cap A'$ a σ -extension of \mathcal{F} ?

Example 6. Let us consider again $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ (Fig. 3), and the set $S = \{a, b\}$. Since S is not conflict-free regarding the certain attack relation R , it is not conflict-free in any completion of \mathcal{I} , thus it cannot be an extension for any semantics considered in this paper. The pair (\mathcal{I}, S) is thus a “NO” instance for σ -IncPV and σ -IncNV.

Now, consider $S' = \{a, c, d, e\}$. S' is a σ -extension ($\sigma \in \{co, pr, stb, gr\}$) for some completions of \mathcal{I} (e.g. \mathcal{F}_8 , see Fig. 4), so (\mathcal{I}, S') is a “YES” instance for σ -IncPV. It is, however, a “NO” instance for σ -IncNV, because it is not conflict-free in completions where (a, c) exists (e.g. \mathcal{F}_{12}).

Finally, we show that there is no “YES” instance with \mathcal{I} for σ -IncNV (for σ one of the semantics considered here). Towards a contradiction, suppose that there is some set S'' such that $S'' \cap A'_i$ is a σ -extension of $\mathcal{F}_i = \langle A'_i, R'_i \rangle$ for each completion $\mathcal{F}_i \in \text{comp}(\mathcal{I})$. S'' must contain c, d and e since they are unattacked in e.g. \mathcal{F}_{12} , they appear in every σ -extension of it. For similar reasons, S'' must contain b (which is unattacked in e.g. \mathcal{F}_4 , thus it belongs to every extension of it). Already we face a contradiction: if $b \in S''$, then $S'' \notin \sigma(\mathcal{F}_{12})$ because b is attacked by the (unattacked) arguments d and e in \mathcal{F}_{12} .

Further research on extension verification has been conducted. In [50], a set of arguments for which the answer to σ -IncPV (respectively σ -IncNV) is “YES” is called a possible (respectively necessary) i -extension. The authors identify some issues with this definition (for instance, a set of arguments could be identified as a possible i -extension even if it is not conflict-free). To remedy this issue, they define i^* -extensions, and the corresponding verification problems:

σ -IncPV* Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $S \subseteq A \cup A^?$, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ such that S is a σ -extension of \mathcal{F} ?

σ -IncNV* Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $S \subseteq A \cup A^?$, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, is S a σ -extension of \mathcal{F} ?

Example 7. Observe that in $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ (Fig. 3), $S = \{b, c, e\}$ is a possible σ i -extension ($\sigma \in \{co, pr, stb, gr\}$), i.e. (\mathcal{I}, S) is a “YES” instance for σ -IncPV. For instance, in $\mathcal{F}_1 = \langle A'_1, R'_1 \rangle$ (Fig. 4), $S \cap A'_1 = \{b, c\}$ is a σ -extension. It is counterintuitive, since S is not conflict-free in completions where $S \subseteq A'_i$. On the contrary, (\mathcal{I}, S) is a “NO” instance for σ -IncPV*, since there is no completion such that S is an extension of it.

We refer the interested reader to [50] for a more detailed discussion of the difference between i -extensions and i^* -extensions.

²See [16,19] for preliminary results regarding verification in Att-IAFs and Arg-IAFs.

Argument acceptability. Then, the (possible and necessary) variants of credulous and skeptical acceptability are studied in [15,17]:

- σ -PCA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ such that $a \in \bigcup_{E \in \sigma(\mathcal{F})} E$?
- σ -NCA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, is $a \in \bigcup_{E \in \sigma(\mathcal{F})} E$?
- σ -PSA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ such that $a \in \bigcap_{E \in \sigma(\mathcal{F})} E$?
- σ -NSA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, is $a \in \bigcap_{E \in \sigma(\mathcal{F})} E$?

Example 8. We continue the previous example, with $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ from Fig. 3. Notice that, since all the completions are acyclic graphs, they all have one extension which is complete, preferred, stable and grounded at the same time (this result is known since Dung's seminal paper [44]). (\mathcal{I}, a) is a "YES" instance for σ -PCA and σ -PSA for all considered semantics, since a belongs to the (unique) extension of some completions of \mathcal{I} (e.g. \mathcal{F}_3), but it is a "NO" instance for σ -NCA and σ -NSA because a is not accepted in each completion (see e.g. \mathcal{F}_1). On the contrary, (\mathcal{I}, c) is a "YES" instance for σ -NCA and σ -NSA, since c belongs to the (unique) extension of every completion.

Non-empty existence. The variants of non-empty existence have been defined in [77,78]:

- σ -PNE Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ with a non-empty σ -extension?
- σ -NNE Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, does \mathcal{F} have a non-empty σ -extension?

Example 9. The IAF \mathcal{I} from Fig. 3 is a "YES" instance for both σ -PNE and σ -NNE, since all the completions possess (exactly) one non-empty extension for all the semantics under consideration.

Actually, [77] also mentions the "existence problem", where the answer is true if there is an extension (empty or not). For most semantics, the answer to this question is trivially "YES" for any AF (thus, for any completion of any IAF). The only case where this problem is not trivial is the stable semantics, but then it coincides with the non-empty existence problem, so we choose not to give more details about this (less demanding) version.

The complexity of these problems under various classical semantics is summarized in Table 4.

Existence and skeptical acceptability. Like in the case of classical AFs, the ExSA problem (or more precisely its adaptations to IAFs) has only been studied for $\sigma \in \{cf_{\neq \emptyset}, adm_{\neq \emptyset}, stb\}$. The formal definition

Table 4
Complexity of reasoning with IAFs under $\sigma \in \{adm, stb, co, gr, pr\}$ [15,18,50,77]. C-c means C-complete

σ	σ -IncPV	σ -IncNV	σ -IncPV*	σ -IncNV*	σ -PCA	σ -NCA	σ -PSA	σ -NSA	σ -PNE	σ -NNE
<i>adm</i>	NP-c	P	P	P	NP-c	Π_2^P -c	trivial	trivial	NP-c	Π_2^P -c
<i>stb</i>	NP-c	P	P	P	NP-c	Π_2^P -c	Σ_2^P -c	coNP-c	NP-c	Π_2^P -c
<i>co</i>	NP-c	P	P	P	NP-c	Π_2^P -c	NP-c	coNP-c	NP-c	Π_2^P -c
<i>gr</i>	NP-c	P	P	P	NP-c	coNP-c	NP-c	coNP-c	P	P
<i>pr</i>	Σ_2^P -c	coNP-c	Σ_2^P -c	coNP-c	NP-c	Π_2^P -c	Σ_3^P -c	Π_2^P -c	NP-c	Π_2^P -c

Table 5
Complexity of σ -PEXSA and σ -NEXSA under $\sigma \in \{cf \neq \emptyset, adm \neq \emptyset, stb\}$ [73]. C-c means C-complete

σ	σ -PEXSA	σ -NEXSA
$cf \neq \emptyset$	in P	in P
$adm \neq \emptyset$	Σ_2^P -c	Π_2^P -c
stb	Σ_2^P -c	Π_2^P -c




Fig. 5. An IAF and its two completions.

is as follows:

- σ -PEXSA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, is there $\mathcal{F} \in \text{comp}(\mathcal{I})$ such that $\sigma(\mathcal{F}) \neq \emptyset$ and $a \in \bigcap_{E \in \sigma(\mathcal{F})} E$?
- σ -NEXSA Given $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ an IAF and $a \in A$, for each $\mathcal{F} \in \text{comp}(\mathcal{I})$, do $\sigma(\mathcal{F}) \neq \emptyset$ and $a \in \bigcap_{E \in \sigma(\mathcal{F})} E$ hold?

Complexity of these problems was also established, results are reported in Table 5.

Example 10. Let $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ be the IAF depicted in Fig. 5(a), with its completions $\text{comp}(\mathcal{I}) = \{\mathcal{F}_1, \mathcal{F}_2\}$ given in Fig. 5(b) and 5(c). \mathcal{F}_1 has no stable extension, while \mathcal{F}_2 has one stable extension $\{b\}$. Thus, (\mathcal{I}, b) is a “YES” instance for stb -PEXSA, and a “NO” instance for stb -NEXSA.

3.3.2. Algorithms

A computational method based on Boolean satisfiability has been proposed for acceptance problems in IAFs [73]. Some problems are at the first level of the polynomial hierarchy, and thus can be translated into instances of the SAT or UNSAT problems. Then, a CEGAR (CounterExample Guided Abstraction Refinement [31]) algorithm based on a SAT oracle is proposed for second level complete problems. Let us briefly describe the Boolean encodings, based on standard translation of semantics by [21].

Given an IAF $\mathcal{I} = \langle A, A^?, R, R^? \rangle$, define Boolean variables x_a, y_a for each $a \in A \cup A^?$, and $r_{a,b}$ for all $(a, b) \in R \cup R^?$, with the following meaning:

- $y_a = 1$ if and only if $a \in A'$,
- $r_{a,b} = 1$ if and only if $(a, b) \in R'$,
- and $x_a = 1$ if and only if $\exists E \in \sigma(\mathcal{F}')$ such that $a \in E$,

where $\mathcal{F}' = \langle A', R' \rangle$ is a completion of \mathcal{I} associated with the y_a and $r_{a,b}$ variables, and σ is a semantics. The structural information about completions is encoded by

$$\phi_{\mathcal{I}} = \bigwedge_{a \in A} y_a \wedge \bigwedge_{(a,b) \in R} r_{a,b} \wedge \bigwedge_{a \in A^?} \left(\neg y_a \rightarrow \left(\neg x_a \wedge \bigwedge_{(a,b) \in R^?} \neg r_{a,b} \wedge \bigwedge_{(b,a) \in R^?} \neg r_{b,a} \right) \right)$$

that expresses the fact that certain arguments and attacks must be true, and uncertain arguments that do not appear in a completion cannot be accepted or be involved in attacks. In the remainder of this part,

we focus on the stable semantics for a matter of exemplification, but the original paper also describes a Boolean encoding for admissibility. First, we describe how conflict-freeness is encoded:

$$\phi_{cf} = \bigwedge_{(a,b) \in R \cup R^?} ((y_a \wedge y_b \wedge r_{a,b}) \rightarrow (\neg x_a \vee \neg x_b))$$

expresses the fact that two arguments cannot be jointly accepted (with respect to a completion) if there is an attack between them (in the said completion). Then, the stable semantics is encoded in

$$\phi_{stb} = \phi_{cf} \wedge \bigwedge_{a \in A \cup A^?} ((\neg x_a \wedge y_a) \rightarrow z_a)$$

where z_a is a newly introduced variable, for each $a \in A \cup A^?$, with

$$z_a \leftrightarrow \bigvee_{(b,a) \in R \cup R^?} (x_b \wedge y_b \wedge r_{b,a}).$$

In words, an argument a that appears in a completion (y_a) without being accepted ($\neg x_a$) must be attacked by some accepted argument (z_a). Then, the NP-complete problem *stb*-PCA can be solved by checking whether the formula $\phi_{stb}^{PCA}(\mathcal{I}, a) = \phi_?(\mathcal{I}) \wedge \phi_{stb}(\mathcal{I}) \wedge x_a$ is satisfiable, while the coNP-complete problem *stb*-NSA corresponds to checking whether $\phi_{stb}^{NSA} = \phi_?(\mathcal{I}) \wedge \phi_{stb}(\mathcal{I}) \wedge \neg x_a$ is unsatisfiable. We refer the interested reader to [15,73] for more details about the encodings and the CEGAR algorithm for the second level of the polynomial hierarchy.

3.4. Direct semantics for PAFs and IAFs

Besides the possible and necessary reasoning modes described previously, [29] defines semantics for Partial AFs (recall that it corresponds to Att-IAFs, *i.e.* IAFs with $A^? = \emptyset$), based on alternative definitions of the basic notions of conflict-freeness and acceptability that underly Dung's semantics. Here we call these semantics "direct semantics", opposed to the completion-based reasoning presented before. In the following, we adapt the definitions of [29] to be consistent with the notations used in this paper.

Definition 7 (Conflict-freeness in PAFs). Given a PAF $\mathcal{P} = \langle A, R, R^? \rangle$, a set of arguments $S \subseteq A$ is

- R -conflict-free if and only if $\nexists a, b \in S$ such that $(a, b) \in R$;
- R - $R^?$ -conflict-free if and only if $\nexists a, b \in S$ such that $(a, b) \in R \cup R^?$.

Intuitively, R -conflict-freeness considers that the incompatibility between a and b , when $(a, b) \in R^?$, is not severe enough to prevent a and b of being accepted together. Thus, uncertain attacks can be ignored for identifying R -conflict-freeness. On the contrary, R - $R^?$ -conflict-freeness excludes all conflicts (both certain and uncertain) from acceptable sets.

Example 11. Consider $\mathcal{P} = \langle A, R, R^? \rangle$ from Fig. 6. The set of arguments $S = \{a_2, a_3\}$. is R -conflict-free, but not R - $R^?$ -conflict-free. Then, $S' = \{a_1, a_3\}$ is R - $R^?$ -conflict-free (and naturally, R -conflict-free as well).




Fig. 6. A PAF \mathcal{P} .

In standard AFs, the notion of admissibility combines conflict-freeness and self-defense [44]. More precisely, Dung defines an argument a as *acceptable with respect to a set of arguments* S if each attacker of a is attacked by an element in S . This notion is also extended to PAFs.

Definition 8 (Acceptability in PAFs). Given a PAF $\mathcal{P} = \langle A, R, R^? \rangle$ and a set of arguments $S \subseteq A$, an argument $a \in A$ is

- R -acceptable with respect to S if $\forall b \in A$ such that $(b, a) \in R, \exists c \in S$ such that $(c, b) \in R$;
- R - $R^?$ -acceptable with respect to S if $\forall b \in A$ such that $(b, a) \in R \cup R^?, \exists c \in S$ such that $(c, b) \in R$;

Similarly to R -conflict-freeness, R -acceptability considers uncertain attacks as “not serious enough” for requiring to be counterattacked. On the contrary, R - $R^?$ -acceptability implies that an argument is defended against all its attacks, not only the certain ones. Notice that a quite cautious view is adopted here, since only certain attacks can be used to defend an argument (even against uncertain attacks).

Example 12. Consider again $\mathcal{P} = \langle A, R, R^? \rangle$ from Fig. 6. The argument a_3 is R -acceptable with respect to $S = \{a_3\}$: the only attacker of a_3 (with respect to R) is a_4 , and $(a_3, a_4) \in R$. Since a_2 is a “weaker” attacker (*i.e.* $(a_2, a_3) \in R^?$ instead of R), there is no need for a_3 to defend itself against a_2 . But then, a_3 is not R - $R^?$ -acceptable with respect to S , since S does not defend a_3 against a_2 . On the contrary, a_3 is R - $R^?$ -acceptable with respect to $S' = \{a_1, a_3\}$. Notice that it would not be the case if (a_1, a_2) was in $R^?$ instead of R .

Then, different versions of admissibility can be defined by combining a notion of conflict-freeness with a notion of acceptability. More precisely:

Definition 9 (Admissibility in PAFs). Given a PAF $\mathcal{P} = \langle A, R, R^? \rangle$, a set of arguments $S \subseteq A$ is

- admissible if and only if S is R -conflict-free, and $\forall a \in S, a$ is R -acceptable with respect to S ;
- R -admissible if and only if S is R - $R^?$ -conflict-free, and $\forall a \in S, a$ is R -acceptable with respect to S ;
- R - $R^?$ -admissible if and only if S is R - $R^?$ -conflict-free, and $\forall a \in S, a$ is R - $R^?$ -acceptable with respect to S .

One may wonder why there is no notion of admissibility characterized by R -conflict-freeness and R - $R^?$ -acceptability. It is proven that such a set would actually be R - $R^?$ -conflict-free, thus it would be a R - $R^?$ -admissible as well. We use $adm(\mathcal{P}), adm_R(\mathcal{P})$ and $adm_{R,R^?}(\mathcal{P})$ to denote, respectively, the admissible, R -admissible and R - $R^?$ -admissible sets of a PAF \mathcal{P} . The definitions straightforwardly imply that $adm_{R,R^?}(\mathcal{P}) \subseteq adm_R(\mathcal{P}) \subseteq adm(\mathcal{P})$.

Table 6

Complexity of reasoning with weak and strong semantics of IAFs, where $X \in \{S, W\}$ means strong or weak [65]

σ	σ -CA	σ -SA	σ -Ver
adm_X	NP-c	trivial	in P
stb_X	NP-c	coNP-c	in P
co_X	NP-c	in coNP	in P
pr_X	NP-c	Π_2^P -c	coNP-c

Example 13. Let us continue Example 12. We have seen that a_3 is R -acceptable with respect to $S = \{a_3\}$. Moreover, S is R -conflict-free. Thus S is an admissible set of \mathcal{P} . It is also R -admissible, since it is R - $R^?$ -conflict-free. But it is not R - $R^?$ admissible.

The set $S' = \{a_1, a_3\}$ is R - $R^?$ -conflict-free, and all its elements are R - $R^?$ -acceptable with respect to S' . Thus it is R - $R^?$ -acceptable.

Then, [29] focuses on three versions of the preferred semantics, based on the different notions of admissibility defined in Definition 9.

Definition 10 (Preferred Semantics for PAFs). Given a PAF $\mathcal{P} = \langle A, R, R^? \rangle$, the set $S \subseteq A$ is

- a preferred extension if and only if it is a \subseteq -maximal admissible set;
- a R -preferred extension if and only if it is a \subseteq -maximal R -admissible set;
- a R - $R^?$ -preferred extension if and only if it is a \subseteq -maximal R - $R^?$ -admissible set.

The original paper discusses the properties of these semantics (for instance the inclusion links between different semantics, or the relations between the semantics of the PAF and the semantics of the completions) as well as computational issues. Interestingly, contrary to possible and necessary reasoning that can be harder than classical reasoning with AFs (recall Section 3.3.1), reasoning with these new semantics keep the complexity of Dung's AFs. In particular, credulous and skeptical acceptability are (respectively) NP-complete and Π_2^P -complete for the three PAF variants of the preferred semantics.

This approach has been recently generalized to IAFs in [65]. The concepts of conflict-freeness and acceptability are adapted to IAFs under the names weak conflict-freeness (respectively weak defense) corresponding to R -conflict-freeness (respectively R -acceptability) and strong conflict-freeness (respectively strong defense) corresponding to R - $R^?$ -conflict-freeness (respectively R - $R^?$ -acceptability). Then, the three types of admissibility are also defined for IAFs, where "simple" admissibility is renamed weak admissibility, R -admissibility is mixed admissibility, and R - $R^?$ -admissibility is strong admissibility. However, the paper excludes mixed admissibility from further study after proving that it does not satisfy Dung's fundamental lemma, *i.e.* the fact that if an admissible set S defends an argument a , then $S \cup \{a\}$ is admissible as well. Then, [65] adapts the definitions of preferred, complete and stable semantics to weak and strong admissibility, and provides complexity results, reported in Table 6. This paper left open the question of grounded semantics, as well as tight complexity result for skeptical acceptability under (weak and strong) complete semantics.

3.5. Rich IAFs: A closer look to uncertainty of the third kind

Besides uncertain arguments ($A^?$) and uncertain attacks ($R^?$), a third kind of uncertainty can be considered: uncertainty about the direction of attacks. This idea has been introduced in Control AFs [39]

(see Section 4). However, such a new kind of uncertainty can also be applied in the context of IAFs. This leads to the definition of Rich IAFs (RIAFs) [63].

Definition 11 (Rich IAF). A *Rich Incomplete Argumentation Framework* (RIAF) is a tuple $r\mathcal{I} = \langle A, A^?, R, R^?, \leftrightarrow^? \rangle$, where A and $A^?$ are disjoint sets of arguments, and $R, R^?, \leftrightarrow^? \subseteq (A \cup A^?) \times (A \cup A^?)$ are disjoint sets of attacks, such that $\leftrightarrow^?$ is symmetric.

The new relation $\leftrightarrow^?$, borrowed from [39], is a symmetric (uncertain) conflict relation: if $(a, b) \in \leftrightarrow^?$, then we are sure that there is a conflict between a and b , but not of the direction of the attack. This new relation impacts the definition of completions.

Definition 12 (Completion of a RIAF). Given $r\mathcal{I} = \langle A, A^?, R, R^?, \leftrightarrow^? \rangle$, a completion of $r\mathcal{I}$ is $\mathcal{F} = \langle A', R' \rangle$, such that

- $A \subseteq A' \subseteq A \cup A^?$;
- $R_{|A'} \subseteq R' \subseteq R_{|A'} \cup R_{|A'}^? \cup \leftrightarrow_{|A'}^?$;
- if $(a, b) \in \leftrightarrow_{|A'}^?$, then $(a, b) \in R'$ or $(b, a) \in R'$ (or both);

where $R_{|A'} = R \cap (A' \times A')$ (and similarly for $R_{|A'}^?$ and $\leftrightarrow_{|A'}^?$).

Again, we use $\text{comp}(r\mathcal{I})$ to denote the set of completions of a RIAF $r\mathcal{I}$.

Example 14. Figure 7 depicts a RIAF $r\mathcal{I} = \langle A = \{a, b, c, d, e\}, A^? = \{f\}, R = \{(c, a), (d, b), (d, c)\}, R^? = \{(e, a), (f, d)\}, \leftrightarrow^? = \{(a, b), (b, a)\} \rangle$.

We do not give here the full set of completions of $r\mathcal{I}$. Let us focus on one option for each of (e, a) , (f, d) and f , and we only illustrate the three options for (a, b) . These three completions $\mathcal{F}_1, \mathcal{F}_2$ and \mathcal{F}_3 are given at Fig. 8.

Similarly, for each other configuration of (e, a) , (f, d) and f , there are three options for the conflict between a and b , leading to three different completions.

Although this richer framework is strictly more expressive than IAFs (*i.e.* any IAF is a RIAF, but some RIAFs cannot be translated into an equivalent IAF), the complexity of reasoning with RIAFs is the same as the complexity of reasoning with IAFs (see [63] for more details).




Fig. 7. The RIAF $r\mathcal{I}$.




Fig. 8. Three completions of $r\mathcal{I}$.

3.6. Restricting the set of completions

Several approaches have been proposed recently (and concurrently) that allow to restrict the set of completions of an IAF. There are several motivations to do that. As a first natural example of such motivation (borrowed from [53]), consider a situation where two agents Ag_1 and Ag_2 are participating in a debate. Ag_1 knows that Ag_2 may use one of two arguments a_1 and a_2 , such that a_1 (respectively a_2) corresponds to the point of view of right-wing (respectively left-wing) politicians. Ag_1 also knows that Ag_2 is politically oriented, *i.e.* Ag_2 certainly has some preferences between a_1 and a_2 , but she does not know exactly which of right-wing or left-wing corresponds to Ag_2 's orientation. So, the AF of Ag_2 is one of those described at Fig. 9.

However, there is no simple way to represent these AFs with the formalisms presented in previous section, *i.e.* there is no IAF \mathcal{I} with $\text{comp}(\mathcal{I}) = \{\mathcal{F}_1, \mathcal{F}_2\}$. A possible solution is to define an IAF structure that admits a set of completions \mathfrak{F} such that $\{\mathcal{F}_1, \mathcal{F}_2\} \subseteq \mathfrak{F}$, and to add a constraint that excludes from the reasoning all the AFs different from \mathcal{F}_1 or \mathcal{F}_2 .

A similar motivation is given in [66], where the issue of representing any set of AFs as the completions of a (constrained) IAF is related to AF revision [35] and merging [38] whose result might not be representable by a single AF but by a set of AFs instead.

A third motivation is the notion of correlation between arguments [51]. For instance, considering structured argumentation settings [22,23], if a_1 and a_2 are uncertain arguments such that a_1 is a sub-argument of a_2 , then a_2 cannot belong to a completion without a_1 belonging too. Such a constraint can be intuitively represented as $a_2 \Rightarrow a_1$. Different kinds of correlations are proposed by [51], in the context of Arg-IAFs (*i.e.* IAFs with $R^? = \emptyset$).

Definition 13 (Dependency). Given an Arg-IAF $\mathcal{I} = \langle A, A^?, R \rangle$, a dependency δ is

- either $X \Rightarrow Y$,
- or $OP(X)$, $OP \in \{Or, Nand, Choice\}$,

and X, Y are non-empty subsets of $A^?$.

Then, a completion of \mathcal{I} is valid (with respect to δ) under some conditions.

Definition 14 (Valid Completion w.r.t. a Dependency). Given an Arg-IAF $\mathcal{I} = \langle A, A^?, R \rangle$, a completion $\mathcal{F} = \langle A', R' \rangle$ of \mathcal{I} is *valid* with respect to δ if

- $\delta = Or(X)$ and $X \cap A' \neq \emptyset$, *i.e.* at least one argument in X belongs to the completion,
- $\delta = Nand(X)$ and $X \cap A' \subset X$, *i.e.* not all the arguments in X belong to the completion,
- $\delta = Choice(X)$ and $|X \cap A'| = 1$, *i.e.* exactly one argument in X belongs to the completion,
- $\delta = X \Rightarrow Y$ and if $X \subseteq A$ then $Y \cap A' \neq \emptyset$, *i.e.* if all the arguments in X belong to the completion, then at least one argument in Y also does.

An Arg-IAF with dependencies is then a pair made of an Arg-IAF $\mathcal{I} = \langle A, A^?, R \rangle$ and a set of dependencies Δ , and [51] studies the complexity of two decision problems:




Fig. 9. The possible AFs of Ag_2 .

- DSat Given an Arg-IAF with dependencies, is there at least one completion valid with respect to each dependency $\delta \in \Delta$?
- σ -PDVer Given an Arg-IAF with dependencies and a set of argument $S \subseteq A \cup A^?$, is S a σ -extension of some valid completion?

DSat is either trivial or NP-complete, depending on the kinds of dependencies appearing in Δ . σ -PDVer goes from trivial to Σ_2^P -complete depending on the dependencies in Δ and the semantics $\sigma \in \{adm, stb, co, gr, pr\}$.

Arg-IAFs with dependencies have been generalized in [66], in two directions: the Arg-IAF is replaced by a (generalized) IAF, and the set of dependencies is replaced by a propositional formula (over the presence of arguments and attacks). This framework is called Constrained Incomplete Argumentation Framework. The constraint is expressed in a propositional language \mathcal{L}_X built on the Boolean variables $\{\arg_a \mid a \in X\} \cup \{\text{att}_{a_i, a_j} \mid (a_i, a_j) \in X \times X\}$, where X is a set of arguments.

Definition 15 (Constrained IAF). A *Constrained Incomplete Argumentation Framework* (CIAF) is a tuple $c\mathcal{I} = \langle A, A^?, R, R^?, \phi \rangle$, where $\langle A, A^?, R, R^? \rangle$ is an IAF, and $\phi \in \mathcal{L}_{A \cup A^?}$ is a constraint.

Similarly to Arg-IAFs with dependencies, only specific completions of $\langle A, A^?, R, R^? \rangle$ are valid, *i.e.* those which satisfy the constraint ϕ .

Example 15. Let $c\mathcal{I}$ be the CIAF depicted in Fig. 10(a), where $A = \{a_1, a_2\}$, $A^? = \emptyset$, $R = \emptyset$, $R^? = \{(a_1, a_2), (a_2, a_1)\}$ and $\phi = \text{att}_{a_1, a_2} \oplus \text{att}_{a_2, a_1}$. Its completions are $\text{comp}(c\mathcal{I}) = \{\mathcal{F}_1, \mathcal{F}_2\}$ as given respectively in Fig. 10(b) and 10(c). Without the constraint (*i.e.* with $\phi' = \top$), there would be two other completions (one with no attack at all, and one with both attacks (a_1, a_2) , (a_2, a_1)).

The main point of [66] is to prove that any set of AFs can be represented as the completions of a CIAF, and then any set of extensions can be the (union of the sets of) extensions of (the completions of) a CIAF. A similar question has been tackled independently by [53].

Proposition 1 ([53,66]). *Given \mathfrak{F} a set of AFs, there is a CIAF $c\mathcal{I}_{\mathfrak{F}}$ such that $\text{comp}(c\mathcal{I}_{\mathfrak{F}}) = \mathfrak{F}$.*

Proposition 2 ([66]). *Given \mathfrak{S} a set of sets of arguments and $\sigma \in \{co, pr, stb, gr\}$, there is a CIAF $c\mathcal{I}_{\mathfrak{S}}$ such that $\bigcup_{c \in \text{comp}(c\mathcal{I}_{\mathfrak{S}})} \sigma(c) = \mathfrak{S}$.*

Notice that [53] studies CIAFs with another definition, where the graphical structure is completely abstracted away. In this paper, a CIAF is a pair $\langle U, \phi \rangle$ where U is the universe of arguments, and ϕ is a constraint similar to those defined in [66], where the propositional atoms concern the existence of arguments in U and attacks in $U \times U$. This means that such a structure can be mapped with a CIAF $c\mathcal{I} = \langle A = \emptyset, A^? = U, R = \emptyset, R^? = U \times U, \phi \rangle$.




Fig. 10. A CIAF and its completions.

Moreover, the relation between CIAFs and RIAFs (see Section 3.5) is also studied in [53]. It is proven there that any RIAF $r\mathcal{I}$ can be translated into an equivalent CIAF (*i.e.* one with the same completions), while the converse is not true. Since their approach “forgets” the graphical structure of the (C)IAF, one constraint ϕ can be built such that each completion of $r\mathcal{I}$ corresponds to one model of ϕ . Here, we proposed an alternative translation of RIAFs into CIAFs, based on Definition 15.

Definition 16. Given $r\mathcal{I} = \langle A, A^?, R, R^?, \leftrightarrow^? \rangle$ a RIAF, the corresponding CIAF is $c\mathcal{I}_{r\mathcal{I}} = \langle A, A^?, R, R', \phi \rangle$ such that $R' = R^? \cup \leftrightarrow^?$ and $\phi = \bigwedge_{(a,b) \in \leftrightarrow^?} \text{att}_{a,b} \vee \text{att}_{b,a}$.

Applying the definitions of RIAFs and CIAFs, it is easy to see that, for any RIAF $r\mathcal{I}$ such that $c\mathcal{I}_{r\mathcal{I}}$ is its corresponding CIAF, $\text{comp}(r\mathcal{I}) = \text{comp}(c\mathcal{I}_{r\mathcal{I}})$.

4. Control argumentation frameworks

4.1. The model

Now we describe Control Argumentation Frameworks (CAFs) [39], a model that generalizes IAFs in two ways: a new type of uncertainty is added (about the direction of an attack, *i.e.* the same as in RIAFs), and a new kind of arguments and attacks (named control arguments and attacks) is added. A CAF is made of three parts, that represent three types of information from the point of view of an agent:

- (1) the fixed part represent information that cannot be influenced by the agent nor the environment (here, we include other agents as part of the environment), *e.g.* because they are legal information, or arguments that have already been stated in a debate;
- (2) the uncertain part can be seen as information that can be influenced by the environment, *e.g.* arguments that might be used by the other agents in the next steps of the debate;
- (3) finally, the control part represents information that can be influenced by the agent.

Items (1) and (2) are similar to the meaning of certain and uncertain information in IAFs. So, the main conceptual difference between IAFs and CAFs is the control part, that implies some action to be made by the agent, *i.e.* some arguments and attacks should be added by the agent in order to reach a given goal. We will describe these goals in a later subsection.

Let us formally define CAFs:

Definition 17 (Control Argumentation Framework). A Control Argumentation Framework (CAF) is a tuple $\mathcal{C} = \langle F, C, U \rangle$ where F is the fixed part, U is the uncertain part, and C is the control part, with:

- $F = \langle A, R \rangle$, with $R \subseteq (A \cup A^?) \times (A \cup A^?)$,
- $U = \langle A^?, R^? \cup \leftrightarrow^? \rangle$ where $R^?, \leftrightarrow^? \subseteq (A \cup A^?) \times (A \cup A^?)$ and $\leftrightarrow^?$ a symmetric irreflexive relation,
- $C = \langle A_C, R_C \rangle$ where $R_C \subseteq (A_C \times (A \cup A^? \cup A_C)) \cup ((A \cup A^? \cup A_C) \times A_C)$,³

where $A, A^?$ and A_C are disjoint sets of arguments, and $R, R^?, \leftrightarrow^?$ and R_C are disjoint sets of attacks.

$A, A^?, R, R^?$ have the same meaning as in IAFs and $\leftrightarrow^?$ represents the same kind uncertain information as in RIAFs: if $(a, b) \in \leftrightarrow^?$, then either there is actually an attack (a, b) , or (b, a) , or both. This means that there certainly is a conflict between a and b , only the orientation is uncertain. Finally, control

³Equivalently, $R_C \subseteq (A_C \times (A \cup A^?)) \cup ((A \cup A^?) \times A_C) \cup (A_C \times A_C)$.




Fig. 11. A CAF \mathcal{C} .




Fig. 12. A configured CAF \mathcal{C}_{cc} obtained from \mathcal{C} with $cc = \{e, f\}$.

arguments (and the incident control attacks) only exist if the agent chooses to use them. This corresponds to the notion of configuration, formally stated by Definition 19.

Example 16. Figure 11 describes a CAF \mathcal{C} . Bold edges and rectangle nodes represent the control part $C = \{\{e, f, g\}, \{(e, c), (e, d), (e, g), (f, b), (f, d), (g, c), (g, e)\}\}$. Two-heads dashed edges represent $\leftrightarrow^? = \{(a, c), (c, a)\}$. Other nodes and edges have the same meaning as in the case of IAFs.

The aim of an agent using a CAF is to control it with respect to some target set of arguments, *i.e.* make sure (by using some control arguments) that this target belongs to some (or each) extension in some (or each) completion. Let us show how the concept of completion is adapted to CAFs:

Definition 18 (Completion of a CAF). Given a CAF $\mathcal{C} = \langle F, C, U \rangle$, a completion is an AF $\mathcal{F} = \langle A', R' \rangle$ such that

- $A \cup A_C \subseteq A' \subseteq A \cup A_C \cup A^?$,
- $(R \cup R_C)_{|A'} \subseteq R' \subseteq (R \cup R_C \cup R^? \cup \leftrightarrow^?)_{|A'}$, with $(a, b) \in R'$ or $(b, a) \in R'$ for all $(a, b) \in \leftrightarrow^?$.

Finally, a *configuration* of the CAF is the choice of some control arguments (and the incident attacks).

Definition 19 (Control Configuration). Given a CAF $\mathcal{C} = \langle F, C, U \rangle$, a configuration is a set of control arguments $cc \subseteq A_C$. The CAF configured with cc is $\mathcal{C}_{cc} = \langle F, C', U \rangle$ where $C' = \langle cc, R_{C|A'} \rangle$, with $A' = cc \cup A \cup A^?$.

Example 17. We continue the previous example. Figure 12 describes the configured CAF \mathcal{C}_{cc} obtained from \mathcal{C} (Fig. 11) with $cc = \{e, f\}$.

4.2. Complexity of controllability

Controllability is the fact that an agent can configure the CAF in order to make its target set of arguments accepted. In a similar manner to arguments acceptance in IAFs, controllability has four variants: “possible” and “necessary” quantify over the set of completions, while “credulous” and “skeptical” quantify of the set of extensions. The necessary variants were first defined in [39], where it is simply called “controllability”. Then [64] defined the possible versions. Let us formally define the decision problems:

Table 7

Complexity of reasoning with CAFs under $\sigma \in \{adm, stb, co, gr, pr\}$ [39,64,70,72]

σ	σ -PCC	σ -NCC	σ -PSC	σ -NSC	σ -CNE
<i>adm</i>	NP-c	Σ_3^P -c	trivial	trivial	Σ_3^P -c
<i>stb</i>	NP-c	Σ_3^P -c	Σ_2^P -c	Σ_2^P -c	Σ_3^P -c
<i>co</i>	NP-c	Σ_3^P -c	NP-c	Σ_2^P -c	Σ_3^P -c
<i>gr</i>	NP-c	Σ_2^P -c	NP-c	Σ_2^P -c	P
<i>pr</i>	NP-c	Σ_3^P -c	Σ_3^P -c	Σ_3^P -c	Σ_3^P -c

σ -PCC Given $\mathcal{C} = \langle F, C, U \rangle$ a CAF and $T \subseteq A$, is there a configuration $cc \subseteq A_C$ such that T is credulously accepted in some completion of \mathcal{C}_{cc} under σ ?

σ -NCC Given $\mathcal{C} = \langle F, C, U \rangle$ a CAF and $T \subseteq A$, is there a configuration $cc \subseteq A_C$ such that T is credulously accepted in each completion of \mathcal{C}_{cc} under σ ?

σ -PSC Given $\mathcal{C} = \langle F, C, U \rangle$ a CAF and $T \subseteq A$, is there a configuration $cc \subseteq A_C$ such that T is skeptically accepted in some completion of \mathcal{C}_{cc} under σ ?

σ -NSC Given $\mathcal{C} = \langle F, C, U \rangle$ a CAF and $T \subseteq A$, is there a configuration $cc \subseteq A_C$ such that T is skeptically accepted in each completion of \mathcal{C}_{cc} under σ ?

We observe that, for $T = \{a\}$, $\sigma = stb$ and \mathcal{C} from Fig. 11, the answer to all these questions is “YES”: the configured CAF \mathcal{C}_{cc} (Fig. 12) provides an example of a configuration $cc = \{e, f\}$ that guarantees that a is accepted.

Preliminary complexity results for the necessary variants were given in [39], and a more detailed analysis was given in [72]. Complexity results for the possible variants have been given in [64]. While there were no results for *adm*-PCC and *adm*-PSC in [64], we can easily add them: skeptical acceptance is always trivially false since the empty set is admissible in any AF; credulous acceptance with respect to *adm* is always equivalent to credulous acceptance with respect to *pr*. Hence the results presented in Table 7.

Several computational approaches have been defined, based on Boolean satisfiability for the problems at the first level of the polynomial hierarchy, and quantified Boolean formulas (QBFs) [60] or CEGAR [31] for higher complexity problems. The intuition behind these approaches is similar to the one for IAFs, and again logical encodings are based on [21].

Finally, recent work [70] has adapted the work on non-emptiness from IAFs [77,78] to CAFs:

σ -CNE Given $\mathcal{C} = \langle F, C, U \rangle$ a CAF, is there a configuration $cc \subseteq A_C$ such that for each completion $\mathcal{F} \in \text{comp}(\mathcal{C}_{cc})$, $\exists S \in \sigma(\mathcal{F})$ with $S \neq \emptyset$?

This problem of *Control Non-Emptiness* is (for most semantics) harder than the equivalent problems for IAFs. More precisely, the problem is Σ_3^P -complete under all the semantics considered here, except the grounded semantics where it is polynomial (see Table 7).⁴ [70] also mentions the fact that a “possible” variant of control non-emptiness (inspired by possible controllability [64]) would be NP-complete for all $\sigma \in \{adm, co, pr, stb\}$.

⁴It is also proven in [70] that *cf*-CNE is polynomial.

5. Application scenarios

Now we describe three scenarios where IAFs or CAFs are useful. First, the aggregation of several AFs into a single one is the context where IAFs (or more precisely PAFs) were first proposed. Then, we present stability, a concept initially introduced for structured argumentation. It has been proven that adapting this notion to abstract argumentation is strongly connected to reasoning with IAFs. Finally, CAFs have shown some interest in the context of automated negotiation.

5.1. AF merging

Partial AFs [33] were defined as a tool in a process of AF merging. The goal is to obtain a global view of the knowledge of a group of agents Ag_1, \dots, Ag_n , where each agent has her personal AF $\mathcal{F}_i = \langle A_i, R_i \rangle$, where both A_i and R_i can differ from an agent to another. The AF of an agent is expanded into a PAF that represents, somehow, the agent's opinion about the group knowledge. Different expansion policies can be used. Here we only describe the *consensual expansion*:

Definition 20 (Consensual Expansion of an AF). Given $\mathcal{F}_1 = \langle A_1, R_1 \rangle, \dots, \mathcal{F}_n = \langle A_n, R_n \rangle$ n AFs corresponding to the knowledge of n agents, the consensual expansion of $\mathcal{F}_i = \langle A_i, R_i \rangle$ is $\mathcal{P}_i = \langle A, R, R^? \rangle$, defined by:

- $A = \bigcup_{j=1}^n A_j$,
- $\forall a, b \in A_i, (a, b) \in R$ iff $(a, b) \in R_i$,
- $\forall a, b \in A$ such that $a \notin A_i$ or $b \notin A_i$, define $\mathcal{A} = \{\mathcal{F}_j \mid a, b \in A_j\}$.
 - * If $\forall \mathcal{F}_j \in \mathcal{A}, (a, b) \in R_j$, then $(a, b) \in R$.
 - * If $\exists \mathcal{F}_j, \mathcal{F}_k \in \mathcal{A}$ such that $(a, b) \in R_j$ and $(a, b) \notin R_k$, then $(a, b) \in R^?$.
 - * Otherwise, (a, b) is included neither in R nor $R^?$.

All the arguments that appear in one of the AFs also appear in agent Ag_i 's PAF. If two arguments a and b both appear in her initial AF, then the existence (or non-existence) and direction of attacks between a and b is the same as in \mathcal{F}_i . Then, for the arguments that agent Ag_i does not initially know, she somehow trusts the other agents. More precisely, if all the agents that know both a and b agree on the fact that there is an attack from a to b , then there is a certain attack (a, b) in the PAF. If they disagree (*i.e.* some of them consider that there is an attack (a, b) , and some other consider that there is no such attack), then (a, b) becomes an uncertain attack in the PAF. Finally, if they all agree that there is no attack (a, b) , then there is no attack (a, b) in the PAF (neither certain nor uncertain).

Example 18. Figure 13 describes three AFs \mathcal{F}_i and their respective expansion \mathcal{P}_i , for $i \in \{1, 2, 3\}$, corresponding to three agents. Let us focus on \mathcal{F}_1 . The relations between a and b correspond to those in \mathcal{F}_1 (*i.e.* a attacks b , but not the other way around), since Ag_1 initially knows these arguments. Arguments c and d are added, and the other agents AFs are used to define the attack relation concerning these arguments. Only one agent (Ag_3) knows c and d , so the relation between c and d in \mathcal{P}_1 is the same as in \mathcal{F}_3 : c attacks d . For b and c , the situation is different. Two agents (Ag_2 and Ag_3) know both b and c , but they disagree and the relation between them: the attack (b, c) only exists in \mathcal{F}_2 . So (b, c) is included in the uncertain attack relation of \mathcal{P}_1 .

The expansion of AFs into PAFs allows to compare each agent's knowledge with the same set of arguments. The result of the merging is then a set of AFs, built on A , that minimize the (aggregated) distance with the set of PAFs that correspond to the expansions of $\mathcal{F}_1, \dots, \mathcal{F}_n$.

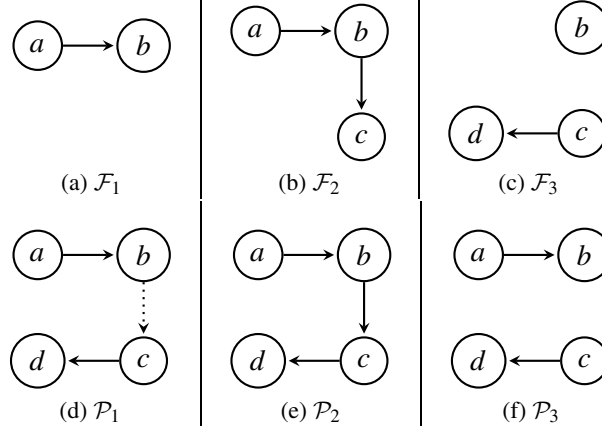


Fig. 13. Three AFs \mathcal{F}_i and their expansions \mathcal{P}_i , $i \in \{1, 2, 3\}$.

Definition 21 (Merging Operator). Given $\mathcal{F}_1 = \langle A_1, R_1 \rangle, \dots, \mathcal{F}_n = \langle A_n, R_n \rangle$ n AFs corresponding to the knowledge of n agents, the merging of $\mathcal{F}_1, \dots, \mathcal{F}_n$ is defined by:

$$\Delta_d^\oplus(\mathcal{F}_1, \dots, \mathcal{F}_n, \mathcal{P}_1, \dots, \mathcal{P}_n) = \left\{ \mathcal{F}^* = \langle A^*, R^* \rangle \mid \mathcal{F}^* \text{ minimizes } \bigoplus_{i=1}^n d(\mathcal{F}^*, \mathcal{P}_i) \right\}$$

where d is a distance between PAFs, \oplus is an aggregation function, and $A^* = \bigcup_{i=1}^n A_i$.

The operator \oplus can be any classical aggregation function, *e.g.* the sum, the max or the leximax. Here, for a matter of simplification we have only described the consensual expansion, but each agent can have her own expansion policy, which means that all the \mathcal{P}_i do not need to be built with the same method. Finally, let us give an example of distance between PAFs.

Definition 22 (Edit Distance). Given two PAFs $\mathcal{P}_1 = \langle A, R_1, R_1^? \rangle$ and $\mathcal{P}_2 = \langle A, R_2, R_2^? \rangle$, the *edit distance over a and b* (with $a, b \in A$) is defined by:

$$de_{a,b}(\mathcal{P}_1, \mathcal{P}_2) = \begin{cases} 0 & \text{if and only if } (a, b) \in R_1 \cap R_2, \text{ or } (a, b) \in R_1^? \cap R_2^?, \text{ or } (a, b) \in N_1 \cap N_2 \\ 1 & \text{if and only if } (a, b) \in R_1 \cap N_2, \text{ or } (a, b) \in N_1 \cap R_2 \\ 0.5 & \text{otherwise} \end{cases}$$

where $N_i = (A \times A) \setminus (R_i \cup R_i^?)$. Then, the *edit distance between \mathcal{P}_1 and \mathcal{P}_2* is

$$de(\mathcal{P}_1, \mathcal{P}_2) = \sum_{a,b \in A} de_{a,b}(\mathcal{P}_1, \mathcal{P}_2)$$

The distance between two PAFs is defined as the sum of the scores of all the possible pairs of arguments, where the score of a pair is 0 if both PAFs agree on the nature of the pair (certain attack, uncertain attack, or absence of attack), 1 if they strongly disagree (certain attack for one PAF, absence of attack for the other one) and 0.5 otherwise (certain attack or absence of attack for one PAF, and uncertain attack for the other one). Since AFs are specific PAFs (with $R^? = \emptyset$), the distance de can be applied on AFs.




Fig. 14. Result of the merging.

Table 8

Distances between AFs and PAFs				
\mathcal{F}_i^*	\mathcal{P}_1	\mathcal{P}_2	\mathcal{P}_3	\sum
\mathcal{F}_1^*	0.5	0	1	1.5
\mathcal{F}_3^*	0.5	1	0	1.5

Example 19. Now we illustrate the result of the merging operator Δ_{de}^Σ , with $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ from Example 18. Again, for a matter of simplification, we consider that all the agents use the consensual expansion (Definition 20). So, the AFs and their respective expansions are those from Fig. 13.

We have $\Delta_{de}^\Sigma(\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3) = \{\mathcal{F}_1^*, \mathcal{F}_2^*\}$, with \mathcal{F}_1^* and \mathcal{F}_2^* from Fig. 14. We detail the computation of $de(\mathcal{F}_1^*, \mathcal{P}_1)$. For $(x, y) \in \{(a, b), (c, d)\}$, both \mathcal{F}_1^* and \mathcal{P}_1 agree on the fact that (x, y) certainly exist, so $de_{x,y}(\mathcal{F}_1^*, \mathcal{P}_1) = 0$. For (b, c) , we obtain $de_{b,c}(\mathcal{F}_1^*, \mathcal{P}_1) = 0.5$ since it is an uncertain attack in \mathcal{P}_1 , and a certain attack in \mathcal{F}_1^* . Finally, for any other (x, y) , both \mathcal{F}_1^* and \mathcal{P}_1 agree on the absence of attack, so $de_{x,y}(\mathcal{F}_1^*, \mathcal{P}_1) = 0$. All the distances, as well as their aggregation with \sum , are presented in Table 8.

5.2. Stability

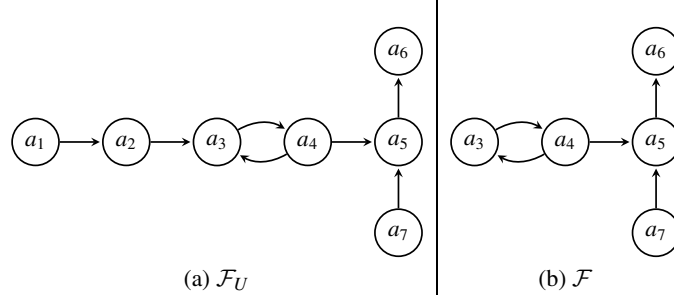
Stability was defined in a context of structured argumentation: it consists in determining whether a given literal will keep the same acceptance status with respect to the grounded semantics, whatever the argumentation setting evolutions [74,81]. It has been applied in the design of an AI agent that helps the Dutch police in investigations about Internet trade frauds. The concept was then adapted to abstract argumentation (and generalized to other semantics) in [67], where the link with IAFs is made.

The hypothesis is made that there is an ‘‘argumentation universe’’ $\mathcal{F}_U = \langle A_U, R_U \rangle$ that describes all the possible arguments and the attacks between them. Any valid AF is then $\mathcal{F} = \langle A, R \rangle$ with $A \subseteq A_U$ and $R = R_U|_A$. Now, for determining whether an AF is stable, we need a way to characterize its possible evolution. It is made with the notion of *future AF*.

Definition 23 (Future Argumentation Framework). Given an AF $\mathcal{F} = \langle A, R \rangle$, the *future AFs* are $F(\mathcal{F}) = \{\mathcal{F}' = \langle A', R' \rangle \mid A \subseteq A', R' = R_U|_{A'}\}$.

Since it is assumed that any valid AF is a subgraph of \mathcal{F}_U , the number of future AFs is finite; precisely $|F(\mathcal{F})| = 2^k$, with $k = |A_U \setminus A|$. Stability with respect to a given argument $a \in A$ consists in analyzing all the future AFs to determine whether the acceptability status of a is the same in all the future AFs. Note that $\mathcal{F} \in F(\mathcal{F})$.

Definition 24 (Stability). An AF $\mathcal{F} = \langle A, R \rangle$ is said to be credulously (respectively skeptically) σ -stable with respect to some $a \in A$ if and only if

Fig. 15. The argumentation universe \mathcal{F}_U and a possible AF \mathcal{F} .

- either $\forall \mathcal{F}' \in F(\mathcal{F})$, a is credulously (respectively skeptically) accepted with respect to σ ;
- or $\forall \mathcal{F}' \in F(\mathcal{F})$, a is not credulously (respectively skeptically) accepted with respect to σ .

Example 20. Let us borrow the example from [67]. Consider the argumentation universe \mathcal{F}_U and the AF \mathcal{F} , both depicted in Fig. 15. The argument a_3 is not credulously σ -stable for $\sigma = stb$, since it is credulously accepted in \mathcal{F} , but not in the future AF where a_2 is added. On the contrary, it is skeptically σ -stable since it is not skeptically accepted in \mathcal{F} , nor in any future AF.

a_6 is skeptically σ -stable as well, but for another reason: indeed we observe that in \mathcal{F} (and in any future AF), a_6 is defended by the (unattacked) argument a_7 , thus it belongs to every extension.

It has been proven that deciding stability can be done through Arg-IAFs, *i.e.* $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ with $R^? = \emptyset$. For a matter of simplicity, we omit $R^?$ in the rest of the section. Let us show how an AF can be associated with an IAF such that the future AFs correspond to the completions of the IAF.

Definition 25 (IAF Corresponding to an AF). Given an AF $\mathcal{F} = \langle A, R \rangle$, the corresponding IAF is $\mathcal{I}_{\mathcal{F}} = \langle A, A \setminus A_U, R_U \rangle$.

It means that all the arguments that appear in \mathcal{F} are certain, and all the other ones are uncertain. Finally, all the attacks are certain. The correspondance between $F(\mathcal{F})$ and $\text{comp}(\mathcal{I}_{\mathcal{F}})$ is straightforward. It implies the following result:

Proposition 3. Given an AF $\mathcal{F} = \langle A, R \rangle$, $a \in A$ and σ a semantics, \mathcal{F} is credulously (respectively skeptically) σ -stable with respect to a if and only if

- either a is necessarily credulously (respectively skeptically) accepted in $\mathcal{I}_{\mathcal{F}}$ with respect to σ ;
- or a is not possibly credulously (respectively skeptically) accepted in $\mathcal{I}_{\mathcal{F}}$ with respect to σ .

Stability is still a recent research topic. As illustrated by [74,81], it can be used to improve the communication between a software agent and a human being, stopping the dialogue when the argumentation system is stable with respect to the issue under discussion. More generally, the notion of stability is promising in argument-based dialogue systems, where it can be used either to stop the dialogue when there is no need to continue (thus saving communication resources), or for adapting an agent's strategy. For instance, [67] presents an example of argument-based negotiation where an agent can revise her strategy when she detects that her (current) preferred argument will never be accepted. The correspondance between the notion of stability and IAF-based reasoning means that efficient computational techniques (*e.g.* those from [73] based on Boolean Satisfiability) can be used in all the scenarios where stability plays a role.

5.3. Automated negotiation

Let us present a last application of the frameworks described in this paper. In [40,41], the authors use CAFs in a context of automated negotiation, where they are used to model the agent's knowledge about her opponent. During a negotiation round, the agent uses her own personal theory (*i.e.* an AF) to choose her favorite offer. But then, in order to convince her opponent to accept this offer, the agent selects an argument supporting it in the CAF that models the opponent. In the case where this argument is not directly accepted by the opponent, the agent uses controllability to defend her argument with control arguments. The technical details of the approach are out of the scope of this paper, we choose to focus on an example of negotiation dialogue that allows to illustrate its main features.

Example 21. Let $\mathcal{F}_1, \mathcal{F}_2$ be the AFs that represent the personal knowledge of the negotiating agents Ag_1, Ag_2 , and $\mathcal{C}_{1,2}, \mathcal{C}_{2,1}$ the CAFs that represent their knowledge about the opponent (*i.e.* $\mathcal{C}_{X,Y}$ represents the knowledge of agent X about agent Y). These AFs and CAFs are shown at Fig. 16. Suppose that reasoning is based on credulous acceptance with respect to $\sigma = stb$. Before the negotiation, Ag_1 prefers the offer o_1 (supported by the accepted argument x in \mathcal{F}_1), and Ag_2 prefers o_2 (supported by the accepted argument v in \mathcal{F}_2). Ag_1 starts the negotiation. The dialogue is made of the following steps.

- (1) *Best offer selection.* Agent Ag_1 need to determine which offer among o_1 and o_2 she prefers, using her personal theory. As already mentioned, o_1 is supported by an accepted argument (x) in \mathcal{F}_1 . On the contrary, o_2 is only supported by a rejected argument (z). So o_1 is preferred to o_2 .
- (2) *Supporting argument selection.* For convincing her opponent to accept o_1 , agent Ag_1 will use her knowledge about Ag_2 . More precisely, she needs to make the argument y accepted in Ag_2 's theory (since y is the only argument in $\mathcal{C}_{1,2}$ that supports o_1).
- (3) *Bidding strategy.* Without using control arguments, it is impossible to convince agent Ag_2 to accept y , since there would be completions where y is not defended against b or e . Thus the agent

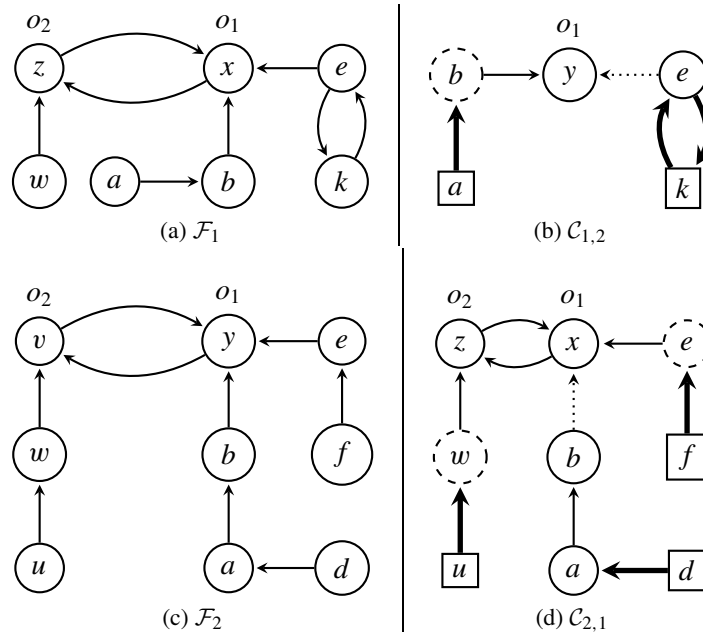


Fig. 16. Agents Ag_1, Ag_2 personal theories (left) and opponent modeling (right).

searches for a control configuration, and finds that $cc = \{a, k\}$ ensures the acceptability of y in every completion of $\mathcal{C}_{1,2}$. Thus Ag_1 proposes “offer o_1 , supported by y , that is accepted because of $\{a, k\}$ ”.

- (4) *Acceptance strategy*. When receiving an offer, agent Ag_2 updates its theory and CAF accordingly. Here, she already knows the arguments y and a . She adds k as well as the attacks between e and k . However, after this update, y remains rejected because b is accepted (it is defended by d). She rejects the offer, and communicates the reasons to agent Ag_1 .
- (5) *Update and new supporting argument selection*. Agent Ag_1 updates her CAF with the information given by Ag_2 . The argument y cannot be used any longer. The agent searches for another supporting argument. Since there is none, Ag_1 's round is over and the token goes to agent Ag_2 .
- (6) *Best offer selection*. Now, agent Ag_2 needs to choose an offer. Since o_2 is supported by the accepted argument v , and o_1 is only supported by a rejected argument (y), she chooses o_2 .
- (7) *Supporting argument selection*. Using her CAF, she determines that (only) z supports the offer o_2 .
- (8) *Bidding strategy*. The agent first checks whether z can be accepted without control arguments. It is not the case, since there are completions in which w is present, and then z is not defended. Then, the control configuration $cc = \{u\}$ is sufficient for defending z . So Ag_2 proposes to Ag_1 “offer o_2 , supported by the argument z , that is accepted because of $\{u\}$ ”.
- (9) *Acceptance strategy*. This time, receiving this offer (along with the defending argument u) makes z acceptable in Ag_1 's theory. So she accepts the offer.

The negotiation ends on an agreement.

Experiments [40] have shown that using control arguments in negotiation increases the agreement rate, *i.e.* the percentage of negotiation that reach an agreement between both agents. Even in presence of a high level of uncertainty, the agreement rate is satisfactorily high when the opponent modeling is enough accurate and the number of control arguments is high [41].

6. Related work

The question of qualitative uncertainty has been handled in another way in [11]. In this work, the authors only consider classical AFs, but they introduce undecidedness in the arguments acceptance statuses. This is somehow a dual approach to IAFs, where uncertainty is present in the initial knowledge (presence or absence of arguments and attacks), but not in the outcome of reasoning (arguments acceptance). Several ideas are discussed. First, in some cases, it is not necessary for a particular application to give a precise status to each argument. A possible motivation for that is to avoid unnecessary computations regarding arguments that are not important, or will likely be revised in the future. For dealing with such situation, a labelling-based approach is used. Besides “traditional” labels $\{in, undec, out\}$ [28], a fourth label is introduced that corresponds to situations where the agent “does not care” about [58]. Another similar approach consists in labelling partially the AF, somehow ignoring the arguments that are not considered as interesting. Then, another issue mentioned in [11] is the fact that “ignorance” is not always the same thing, depending on the context. Especially, an agent can be unable to assess the acceptance status of an argument either because she has no information about it, or on the contrary because she has too much (inconsistent) information about it. A preliminary work on epistemological undecidedness, *i.e.* the fact that some piece of information is inherently unknowable, has then been proposed in the context of structured argumentation in [10]. This kind of approach is more related to IAFs, since the uncertainty appears in the input of the reasoning process.

In this paper, we focus on qualitative uncertainty. However, there is a vast literature on quantitative uncertainty in abstract argumentation, namely on probabilistic argumentation. Two families of approaches exist. The epistemic approach (e.g. [56,57,82]) attaches to arguments a degree of belief, *i.e.* the probability (according to the agent) that the argument is acceptable. In that way, probabilistic argumentation can be seen as related to [11]. On the contrary, in the constellation approach (e.g. [52,55,61]), the probability is attached with the arguments and attacks; it is the probability of existence of arguments and attacks. While [55,61] can be seen as quantitative counterparts of IAFs, [52] is the quantitative counterpart of CAFs. The correspondance between IAFs and Probabilistic Argumentation Frameworks (PrAFs) under the constellation approach has been illustrated by [15]. Each completion of a PrAFs is associated with a probability, it is thus possible (given a semantics) to determine the probability that a given set is an extension, or that a given argument is credulously (respectively skeptically) accepted. Then, it is noted that a set of arguments is a possible extension if its probability is strictly greater than 0, and a necessary extension if its probability is 1. Similarly, an argument is necessarily credulously (respectively skeptically) accepted if its probability of credulous (respectively skeptical) acceptability is 1, and it is possibly credulously (respectively skeptically) accepted if its probability of credulous (respectively skeptical) acceptability is strictly greater than 0.

Finally, let us discuss extension enforcement [12–14,36,84]. This operation consists in modifying an argumentation framework $\mathcal{F} = \langle A, R \rangle$ in order to make some set of arguments $S \subseteq A$ (part of) an extension. When considering an IAF $\mathcal{I} = \langle A, A^?, R, R^? \rangle$, determining whether a set of arguments $S \subseteq A$ is (part of) an extension is equivalent to determining whether S can be enforced in \mathcal{F} , by means of the arguments $A^?$ and attacks $R^?$. This is reminiscent of the constraints defined in [83]. Enforcement is also directly related to CAFs, since controllability is the fact that the target set of arguments can be enforced in some (or each) extension of some (or each) completion, by means of control arguments. CAFs provide a richer setting than existing work on extension enforcement, since it includes uncertainty in the argumentation graph, and the control part constrains the possible ways to perform the enforcement.

7. Conclusion

We have described the main approaches that incorporate qualitative uncertainty in abstract argumentation, namely Incomplete Argumentation Frameworks and Control Argumentation Frameworks and variants of these. Figure 17 shows the relation between these models. An arrow $\mathcal{X} \rightarrow \mathcal{Y}$ with a label \mathcal{T} means that any instance of the class \mathcal{X} can be transformed into an instance of the class \mathcal{Y} with the transformation \mathcal{T} . For instance, any AF $\mathcal{F} = \langle A, R \rangle$ is an Arg-IAF $\mathcal{I} = \langle A, A^?, R \rangle$ with $A^? = \emptyset$, and any Arg-IAF is an IAF $\mathcal{I} = \langle A, A^?, R, R^? \rangle$ with $R^? = \emptyset$.

Although these frameworks are still quite recent, they have already shown their interest in various scenarios, and it paves the way to many applications especially in multi-agent systems, where uncertain information is omnipresent, because *e.g.* the cost of communication forces agents to use enthymemes [25], because of trust issues between agents [80], or because of the impossibility to anticipate the environment evolution [59].

Using IAFs, as well as RIAFs, CIAFs and CAFs, in various application scenarios, seems to be a promising line of research. Besides AF merging, stability and negotiation mentioned here, potential applications include argumentation dynamics [43] (especially in logic-based argumentation with uncertain knowledge in the underlying knowledge base [62]), persuasion [27], or risk management [8]. We have identified several other promising research tracks, including the combinations of ideas discussed in this

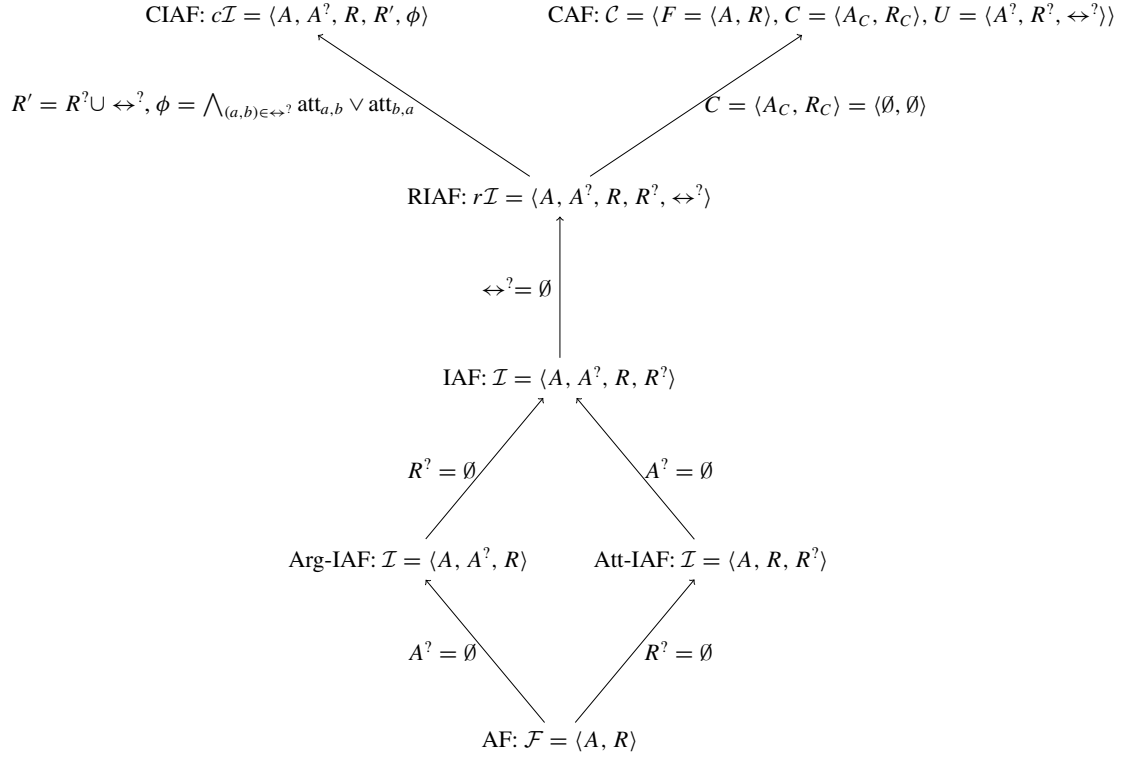


Fig. 17. Relations between Dung's AFs and the frameworks studied in this paper.

paper. For instance, controllability of CAFs could be studied in the context of the extension-based semantics introduced in [65], or with a constrained version of CAFs (using the same kinds of constraints as in [51,53,66]). Similarly, using IAFs instead of PAFs would provide new insights on AF merging [33]. From a more technical point of view, an interesting question consists in determining whether an instance of one of the frameworks presented here can be translated into an equivalent AF (maybe at the cost of additional arguments or attacks, like in the case of Extended AFs [68,69] or Bipolar AFs [30]). This would allow to use existing algorithms dedicated to Dung's AFs for reasoning with these more expressive frameworks. Naturally, the complexity gap for some of these problems means that such translation may not be polynomially doable in some cases. This question should be investigated in depth to identify subclasses of IAFs/CAFs with better computational properties than the general case.

Real world applications of formal argumentation may require a complex combination of different generalizations of Dung's model. A natural question, for instance, is how to combine qualitative uncertainty and bipolar argumentation, *i.e.* how to add (uncertain) supports [3,30] to IAFs. This may also lead to an interesting combination of IAFs (or CAFs) with various frameworks like Weighted AFs [46], Strength-based AFs [76], Value-based AFs [20] or AFs with collective attacks (SETAFs) [71]. The question would then be to determine which arguments are acceptable when the agent is not sure about the weight of some attacks, the strength or the value of some arguments, or the existence of some sub-attack in a collective attack.

Finally, while ranking-based (or gradual) methods for argument acceptability have been a hot topic in the last years (see *e.g.* [1,5,26,37]), there is no such approach for IAFs yet. This would provide valuable

methods for contexts where a more fine-grained reasoning mechanism than a binary acceptability degree is required.

References

- [1] L. Amgoud and J. Ben-Naim, Ranking-based semantics for argumentation frameworks, in: *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, Vol. 8078, Springer, 2013, pp. 134–147. doi:[10.1007/978-3-642-40381-1_11](https://doi.org/10.1007/978-3-642-40381-1_11).
- [2] L. Amgoud and C. Cayrol, A reasoning model based on the production of acceptable arguments, *Ann. Math. Artif. Intell.* **34**(1–3) (2002), 197–215. doi:[10.1023/A:1014490210693](https://doi.org/10.1023/A:1014490210693).
- [3] L. Amgoud, C. Cayrol, M. Lagasque-Schiek and P. Livet, On bipolarity in argumentation frameworks, *Int. J. Intell. Syst.* **23**(10) (2008), 1062–1093. doi:[10.1002/int.20307](https://doi.org/10.1002/int.20307).
- [4] L. Amgoud, Y. Dimopoulos and P. Moraitis, A unified and general framework for argumentation-based negotiation, in: *6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, IFAAMAS, 2007, p. 158.
- [5] L. Amgoud and D. Doder, Gradual semantics accounting for varied-strength attacks, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, Montreal, QC, Canada, May 13–17, 2019, 2019, pp. 1270–1278.
- [6] L. Amgoud and H. Prade, Using arguments for making and explaining decisions, *Artif. Intell.* **173**(3–4) (2009), 413–436. doi:[10.1016/j.artint.2008.11.006](https://doi.org/10.1016/j.artint.2008.11.006).
- [7] S. Arora and B. Barak, *Computational Complexity – A Modern Approach*, Cambridge University Press, 2009, <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264>. ISBN 978-0-521-42426-4.
- [8] T. Aven and O. Renn, Risk management, in: *Risk Management and Governance: Concepts, Guidelines and Applications*, Springer, Berlin, Heidelberg, 2010, pp. 121–158. ISBN 978-3-642-13926-0. doi:[10.1007/978-3-642-13926-0_8](https://doi.org/10.1007/978-3-642-13926-0_8).
- [9] P. Baroni, M. Caminada and M. Giacomin, Abstract argumentation frameworks and their semantics, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, pp. 159–236.
- [10] P. Baroni, M. Giacomin and B. Liao, Dealing with unknowability in formal argumentation, in: *Proceedings of the 1st Workshop on Advances in Argumentation in Artificial Intelligence Co-Located with XVI International Conference of the Italian Association for Artificial Intelligence, AI³@AI*IA 2017*, 2017, pp. 73–78.
- [11] P. Baroni, M. Giacomin and B.S. Liao, I don't care, I don't know ... I know too much! On incompleteness and undecidedness in abstract argumentation, in: *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation – Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, 2015, pp. 265–280. doi:[10.1007/978-3-319-14726-0_18](https://doi.org/10.1007/978-3-319-14726-0_18).
- [12] R. Baumann, What does it take to enforce an argument? Minimal change in abstract argumentation, in: *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, 2012, pp. 127–132.
- [13] R. Baumann and G. Brewka, Expanding argumentation frameworks: Enforcing and monotonicity results, in: *3rd International Conference on Computational Models of Argument (COMMA 2010)*, 2010, pp. 75–86.
- [14] R. Baumann, S. Doutre, J.-G. Mailly and J.P. Wallner, Enforcement in formal argumentation, in: *Handbook of Formal Argumentation*, D. Gabbay, M. Giacomin, G. Simari and M. Thimm, eds, College Publications, 2021, pp. 159–236. To appear in 2021.
- [15] D. Baumeister, M. Järvisalo, D. Neugebauer, A. Niskanen and J. Rothe, Acceptance in incomplete argumentation frameworks, *Artif. Intell.* **295** (2021), 103470. doi:[10.1016/j.artint.2021.103470](https://doi.org/10.1016/j.artint.2021.103470).
- [16] D. Baumeister, D. Neugebauer and J. Rothe, Verification in attack-incomplete argumentation frameworks, in: *4th International Conference on Algorithmic Decision Theory (ADT 2015)*, 2015, pp. 341–358. doi:[10.1007/978-3-319-23114-3_21](https://doi.org/10.1007/978-3-319-23114-3_21).
- [17] D. Baumeister, D. Neugebauer and J. Rothe, Credulous and skeptical acceptance in incomplete argumentation frameworks, in: *7th International Conference on Computational Models of Argument (COMMA 2018)*, 2018, pp. 181–192.
- [18] D. Baumeister, D. Neugebauer, J. Rothe and H. Schadrack, Verification in incomplete argumentation frameworks, *Artif. Intell.* **264** (2018), 1–26. doi:[10.1016/j.artint.2018.08.001](https://doi.org/10.1016/j.artint.2018.08.001).
- [19] D. Baumeister, J. Rothe and H. Schadrack, Verification in argument-incomplete argumentation frameworks, in: *4th International Conference on Algorithmic Decision Theory (ADT 2015)*, 2015, pp. 359–376. doi:[10.1007/978-3-319-23114-3_22](https://doi.org/10.1007/978-3-319-23114-3_22).
- [20] T.J.M. Bench-Capon, Value-based argumentation frameworks, in: *9th International Workshop on Non-monotonic Reasoning (NMR 2002)*, Toulouse, France, April 19–21, Proceedings, S. Benferhat and E. Giunchiglia, eds, 2002, pp. 443–454.
- [21] P. Besnard and S. Doutre, Checking the acceptability of a set of arguments, in: *10th International Workshop on Non-monotonic Reasoning (NMR 2004)*, 2004, pp. 59–64.

- [22] P. Besnard, A.J. García, A. Hunter, S. Modgil, H. Prakken, G.R. Simari and F. Toni, Introduction to structured argumentation, *Argument Comput.* **5**(1) (2014), 1–4. doi:[10.1080/19462166.2013.869764](https://doi.org/10.1080/19462166.2013.869764).
- [23] P. Besnard and A. Hunter, *Elements of Argumentation*, MIT Press, 2008. ISBN 978-0-262-02643-7.
- [24] P. Besnard and A. Hunter, Argumentation based on classical logic, in: *Argumentation in Artificial Intelligence*, G.R. Simari and I. Rahwan, eds, 2009, pp. 133–152. doi:[10.1007/978-0-387-98197-0_7](https://doi.org/10.1007/978-0-387-98197-0_7).
- [25] E. Black and A. Hunter, A relevance-theoretic framework for constructing and deconstructing enthymemes, *J. Log. Comput.* **22**(1) (2012), 55–78. doi:[10.1093/logcom/exp064](https://doi.org/10.1093/logcom/exp064).
- [26] E. Bonzon, J. Delobelle, S. Konieczny and N. Maudet, A comparative study of ranking-based semantics for abstract argumentation, in: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, Phoenix, Arizona, USA, February 12–17, 2016, D. Schuurmans and M.P. Wellman, eds, AAAI Press, 2016, pp. 914–920.
- [27] E. Bonzon, J. Delobelle, S. Konieczny and N. Maudet, A parametrized ranking-based semantics for persuasion, in: *11th International Conference on Scalable Uncertainty Management (SUM 2017)*, Springer, 2017, pp. 237–251. doi:[10.1007/978-3-319-67582-4_17](https://doi.org/10.1007/978-3-319-67582-4_17).
- [28] M. Caminada, On the issue of reinstatement in argumentation, in: *10th European Conference on Logics in Artificial Intelligence (JELIA 2006)*, Springer, 2006, pp. 111–123. doi:[10.1007/11853886_11](https://doi.org/10.1007/11853886_11).
- [29] C. Cayrol, C. Devred and M. 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)*, 2007, pp. 259–270. doi:[10.1007/978-3-540-75256-1_25](https://doi.org/10.1007/978-3-540-75256-1_25).
- [30] C. Cayrol and M. Lagasquie-Schiex, Coalitions of arguments: A tool for handling bipolar argumentation frameworks, *Int. J. Intell. Syst.* **25**(1) (2010), 83–109. doi:[10.1002/int.20389](https://doi.org/10.1002/int.20389).
- [31] E.M. Clarke, O. Grumberg, S. Jha, Y. Lu and H. Veith, Counterexample-guided abstraction refinement for symbolic model checking, *J. ACM* **50**(5) (2003), 752–794. doi:[10.1145/876638.876643](https://doi.org/10.1145/876638.876643).
- [32] S.A. Cook, The complexity of theorem-proving procedures, in: *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing*, Shaker Heights, Ohio, USA, May 3–5, 1971, 1971, pp. 151–158.
- [33] S. Coste-Marquis, C. Devred, S. Konieczny, M. Lagasquie-Schiex and P. Marquis, On the merging of Dung's argumentation systems, *Artif. Intell.* **171**(10–15) (2007), 730–753. doi:[10.1016/j.artint.2007.04.012](https://doi.org/10.1016/j.artint.2007.04.012).
- [34] S. Coste-Marquis, C. Devred and P. Marquis, Symmetric argumentation frameworks, in: *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005*, Barcelona, Spain, July 6–8, 2005, Proceedings, 2005, pp. 317–328.
- [35] S. Coste-Marquis, S. Konieczny, J.-G. Maily and P. Marquis, On the revision of argumentation systems: Minimal change of arguments statuses, in: *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014*, Vienna, Austria, July 20–24, 2014, 2014.
- [36] S. Coste-Marquis, S. Konieczny, J.-G. Maily and P. Marquis, Extension enforcement in abstract argumentation as an optimization problem, in: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, 2015, pp. 2876–2882.
- [37] J. Delobelle, Ranking-based semantics for abstract argumentation (Sémantique à base de classement pour l'argumentation abstraite), PhD thesis, Artois University, Arras, France, 2017. <https://tel.archives-ouvertes.fr/tel-01937279>.
- [38] J. Delobelle, A. Haret, S. Konieczny, J.-G. Maily, J. Rossit and S. Woltran, Merging of abstract argumentation frameworks, in: *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016*, Cape Town, South Africa, April 25–29, 2016, 2016, pp. 33–42.
- [39] Y. Dimopoulos, J.-G. Maily and P. Moraitis, Control argumentation frameworks, in: *32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*, 2018, pp. 4678–4685.
- [40] Y. Dimopoulos, J.-G. Maily and P. Moraitis, Argumentation-based negotiation with incomplete opponent profiles, in: *18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2019)*, 2019, pp. 1252–1260.
- [41] Y. Dimopoulos, J.-G. Maily and P. Moraitis, Arguing and negotiating using incomplete negotiators profiles, *Autonomous Agents and Multi-Agent Systems* **35**(2) (2021).
- [42] Y. Dimopoulos and A. Torres, Graph theoretical structures in logic programs and default theories, *Theor. Comput. Sci.* **170**(1–2) (1996), 209–244. doi:[10.1016/S0304-3975\(96\)80707-9](https://doi.org/10.1016/S0304-3975(96)80707-9).
- [43] S. Doutre and J.-G. Maily, Constraints and changes: A survey of abstract argumentation dynamics, *Argument Comput.* **9**(3) (2018), 223–248. doi:[10.3233/AAC-180425](https://doi.org/10.3233/AAC-180425).
- [44] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artif. Intell.* **77**(2) (1995), 321–358. doi:[10.1016/0004-3702\(94\)00041-X](https://doi.org/10.1016/0004-3702(94)00041-X).
- [45] P.E. Dunne and T.J.M. Bench-Capon, Coherence in finite argument systems, *Artif. Intell.* **141**(1/2) (2002), 187–203. doi:[10.1016/S0004-3702\(02\)00261-8](https://doi.org/10.1016/S0004-3702(02)00261-8).
- [46] P.E. Dunne, A. Hunter, P. McBurney, S. Parsons and M.J. Wooldridge, Weighted argument systems: Basic definitions, algorithms, and complexity results, *Artif. Intell.* **175**(2) (2011), 457–486. doi:[10.1016/j.artint.2010.09.005](https://doi.org/10.1016/j.artint.2010.09.005).
- [47] P.E. Dunne and M.J. Wooldridge, Complexity of abstract argumentation, in: *Argumentation in Artificial Intelligence*, 2009, pp. 85–104. doi:[10.1007/978-0-387-98197-0_5](https://doi.org/10.1007/978-0-387-98197-0_5).

- [48] W. Dvorák and P.E. Dunne, Computational problems in formal argumentation and their complexity, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, pp. 631–688.
- [49] W. Dvorák and S. Woltran, On the intertranslatability of argumentation semantics, *J. Artif. Intell. Res.* **41** (2011), 445–475. doi:[10.1613/jair.3318](https://doi.org/10.1613/jair.3318).
- [50] B. Fazzinga, S. Flesca and F. Furfaro, Revisiting the notion of extension over incomplete abstract argumentation frameworks, in: *29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*, 2020, pp. 1712–1718.
- [51] B. Fazzinga, S. Flesca and F. Furfaro, Reasoning over argument-incomplete AAFs in the presence of correlations, in: *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event, Montreal, Canada, 19–27 August 2021*, Z. Zhou, ed., ijcai.org, 2021, pp. 189–195.
- [52] F. Gaignier, Y. Dimopoulos, J.-G. Mailly and P. Moraitis, Probabilistic control argumentation frameworks, in: *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '21)*, 2021, pp. 519–527.
- [53] A. Herzig and A. Yuste-Ginel, Abstract argumentation with qualitative uncertainty: An analysis in dynamic logic, in: *Proceedings of the Fourth International Conference on Logic and Argumentation (CLAR '21)*, 2021.
- [54] M.J.H. Heule, M. Jarvisalo and M. Suda, SAT competition 2018, *J. Satisf. Boolean Model. Comput.* **11**(1) (2019), 133–154.
- [55] A. Hunter, Probabilistic qualification of attack in abstract argumentation, *Int. J. Approx. Reasoning* **55**(2) (2014), 607–638. doi:[10.1016/j.ijar.2013.09.002](https://doi.org/10.1016/j.ijar.2013.09.002).
- [56] A. Hunter and M. Thimm, On partial information and contradictions in probabilistic abstract argumentation, in: *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR '16)*, 2016, pp. 53–62.
- [57] A. Hunter and M. Thimm, Probabilistic reasoning with abstract argumentation frameworks, *J. Artif. Intell. Res.* **59** (2017), 565–611. doi:[10.1613/jair.5393](https://doi.org/10.1613/jair.5393).
- [58] H. Jakobovits and D. Vermeir, Robust semantics for argumentation frameworks, *J. Log. Comput.* **9**(2) (1999), 215–261. doi:[10.1093/logcom/9.2.215](https://doi.org/10.1093/logcom/9.2.215).
- [59] H. Katsuno and A.O. Mendelzon, On the difference between updating a knowledge base and revising it, in: *2nd International Conference on Principles of Knowledge Representation and Reasoning (KR 1991)*, 1991, pp. 387–394.
- [60] H. Kleine Büning and U. Bubeck, Theory of quantified Boolean formulas, in: *Handbook of Satisfiability*, 2009, pp. 735–760.
- [61] H. Li, N. Oren and T.J. Norman, Probabilistic argumentation frameworks, in: *Proceedings of the First International Workshop on Theory and Applications of Formal Argumentation (TAFa '11)*, 2011, pp. 1–16.
- [62] J.-G. Mailly, Using enthymemes to fill the gap between logical argumentation and revision of abstract argumentation frameworks, *CoRR* **abs/1603.08789** (2016), Accepted at NMR 2016. <http://arxiv.org/abs/1603.08789>.
- [63] J.-G. Mailly, A note on rich incomplete argumentation frameworks, *CoRR* **abs/2009.04869** (2020). <https://arxiv.org/abs/2009.04869>.
- [64] J.-G. Mailly, Possible controllability of control argumentation frameworks, in: *8th International Conference on Computational Models of Argument (COMMA 2020)*, 2020, pp. 283–294.
- [65] J.-G. Mailly, Extension-based semantics for incomplete argumentation frameworks, in: *Proceedings of the Fourth International Conference on Logic and Argumentation (CLAR '21)*, 2021.
- [66] J.-G. Mailly, Constrained incomplete argumentation frameworks, in: *Proceedings of the 16th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU '21)*, J. Vejnárová and N. Wilson, eds, Lecture Notes in Computer Science, Vol. 12897, Springer, 2021, pp. 103–116. doi:[10.1007/978-3-030-86772-0_8](https://doi.org/10.1007/978-3-030-86772-0_8).
- [67] J.-G. Mailly and J. Rossit, Stability in abstract argumentation, *CoRR* **abs/2012.12588** (2020), Accepted at NMR 2020. <https://arxiv.org/abs/2012.12588>.
- [68] S. Modgil, Reasoning about preferences in argumentation frameworks, *Artif. Intell.* **173**(9–10) (2009), 901–934. doi:[10.1016/j.artint.2009.02.001](https://doi.org/10.1016/j.artint.2009.02.001).
- [69] S. Modgil and T.J.M. Bench-Capon, Metalevel argumentation, *J. Log. Comput.* **21**(6) (2011), 959–1003. doi:[10.1093/logcom/exq054](https://doi.org/10.1093/logcom/exq054).
- [70] D. Neugebauer, J. Rothe and K. Skiba, Complexity of nonemptiness in control argumentation frameworks, in: *Proceedings of the 16th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU '21)*, J. Vejnárová and N. Wilson, eds, Lecture Notes in Computer Science, Vol. 12897, Springer, 2021, pp. 117–129. doi:[10.1007/978-3-030-86772-0_9](https://doi.org/10.1007/978-3-030-86772-0_9).
- [71] S.H. Nielsen and S. Parsons, A generalization of Dung's abstract framework for argumentation: Arguing with sets of attacking arguments, in: *Proceedings of the Third International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2006)*, 2006, pp. 54–73.
- [72] A. Niskanen, D. Neugebauer and M. Jarvisalo, Controllability of control argumentation frameworks, in: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, 2020, pp. 1855–1861.

- [73] A. Niskanen, D. Neugebauer, M. Järvisalo and J. Rothe, Deciding acceptance in incomplete argumentation frameworks, in: *34th AAAI Conference on Artificial Intelligence (AAAI 2020)*, 2020, pp. 2942–2949.
- [74] D. Odekerken, A. Borg and F. Bex, Estimating stability for efficient argument-based inquiry, in: *8th International Conference on Computational Models of Argument (COMMA 2020)*, IOS Press, 2020, pp. 307–318.
- [75] C. Papadimitriou and M. Yannakakis, The complexity of facets (and some facets of complexity), *Journal of Computer and System Sciences* **28** (1984), 244–259. doi:[10.1016/0022-0000\(84\)90068-0](https://doi.org/10.1016/0022-0000(84)90068-0).
- [76] J. Rossit, J.-G. Maily, Y. Dimopoulos and P. Moraitis, United we stand: Accruals in strength-based argumentation, *Argument Comput.* **12**(1) (2021), 87–113. doi:[10.3233/AAC-200904](https://doi.org/10.3233/AAC-200904).
- [77] K. Skiba, D. Neugebauer and J. Rothe, Complexity of possible and necessary existence problems in abstract argumentation, in: *24th European Conference on Artificial Intelligence (ECAI 2020)*, 2020, pp. 897–904.
- [78] K. Skiba, D. Neugebauer and J. Rothe, Complexity of nonempty existence problems in incomplete argumentation frameworks, *IEEE Intell. Syst.* **36**(2) (2021), 13–24. doi:[10.1109/MIS.2020.3046782](https://doi.org/10.1109/MIS.2020.3046782).
- [79] L. Stockmeyer, The polynomial-time hierarchy, *Theoretical Computer Science* **3** (1976), 1–22. doi:[10.1016/0304-3975\(76\)90061-X](https://doi.org/10.1016/0304-3975(76)90061-X).
- [80] Y. Tang, K. Cai, P. McBurney, E. Sklar and S. Parsons, Using argumentation to reason about trust and belief, *J. Log. Comput.* **22**(5) (2012), 979–1018. doi:[10.1093/logcom/exr038](https://doi.org/10.1093/logcom/exr038).
- [81] B. Testerink, D. Odekerken and F. Bex, A method for efficient argument-based inquiry, in: *13th International Conference on Flexible Query Answering Systems (FQAS 2019)*, Springer, 2019, pp. 114–125. doi:[10.1007/978-3-030-27629-4_13](https://doi.org/10.1007/978-3-030-27629-4_13).
- [82] M. Thimm, A probabilistic semantics for abstract argumentation, in: *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI '12)*, IOS Press, 2012, pp. 750–755.
- [83] J.P. Wallner, Structural constraints for dynamic operators in abstract argumentation, *Argument Comput.* **11**(1–2) (2020), 151–190. doi:[10.3233/AAC-190471](https://doi.org/10.3233/AAC-190471).
- [84] J.P. Wallner, A. Niskanen and M. Järvisalo, Complexity results and algorithms for extension enforcement in abstract argumentation, *J. Artif. Intell. Res.* **60** (2017), 1–40. doi:[10.1613/jair.5415](https://doi.org/10.1613/jair.5415).