# Quantitative Reasoning over Incomplete Abstract Argumentation Frameworks

**Bettina Fazzinga**[1] , **Sergio Flesca**[2] , **Filippo Furfaro**[2] and **Giuseppina Monterosso**[2]
[1]DiCES - University of Calabria
[2]DIMES - University of Calabria
{bettina.fazzinga, sergio.flesca, filippo.furfaro}@unical.it, giusy.monterosso@dimes.unical.it

## Abstract

We introduce PERCVER and PERCACC, the problems asking for the percentages of the completions of an *incomplete Abstract Argumentation Framework* (iAAF) where a set $S$ is an extension and an argument $a$ is accepted, respectively. These problems give insights into the status of $S$ and $a$ more precise than the "traditional" verification and acceptance tests under the possible and necessary perspectives, that decide if $S$ is an extension and $a$ is accepted in at least one or every completion, respectively. As a first contribution, we study the relationship between the proposed framework and *probabilistic AAFs* (prAAFs) *under the constellations approach* (that, at first sight, seem to be suitable for straightforwardly encoding the quantitative reasoning underlying PERCVER and PERCACC). In this regard, we show that translating an iAAF into an equivalent prAAF requires a heavy computational cost: this backs the study of PERCVER and PERCACC as new distinguished problems. Then, we investigate the complexity of PERCVER and PERCACC, and identify islands of tractability.

## 1 Introduction

Dung's *Abstract Argumentation Framework* (AAF [Dung, 1995]) has proved effective in supporting the reasoning in several scenarios, ranging from the analysis of disputes to process mining tasks [Fazzinga *et al.*, 2022b] and chatbot services [Fazzinga *et al.*, 2022c]. In order to widen the possible applications, several generalizations of AAFs have been proposed to model the uncertainty that may affect arguments and attacks. In fact, in real life disputes, it often happens that the participation of the agent who claims an argument $a$ is not guaranteed (so $a$ is uncertain), or that the existence of an attack $(a, b)$ depends not only from the occurrence of $a$ and $b$, but also on the subjective view of who analyzes the dispute (so $(a, b)$ is uncertain). *Incomplete AAFs* (iAAFs) [Baumeister *et al.*, 2018] are prominent representatives of *qualitative* approaches, where uncertain arguments and attacks can be specified, with no measure of the extent of this uncertainty. Recent proposals [Fazzinga *et al.*, 2021b; Fazzinga *et al.*, 2021a] further extend iAAFs with "dependencies", expressing, for instance, that (the presence of) an argument is alternative to other arguments, that an attack implies another attack, and so on.

The fundamental notion of extension was adapted to iAAFs by taking into account the fact that, while a classical AAF represents a single scenario (in terms of a set of arguments and attacks), an iAAF encodes multiple scenarios (called *completions*), corresponding to the different combinations of presence/absence of the uncertain arguments and uncertain attacks satisfying the dependencies. Thus, based on the notion of completion, *possible and necessary i*$^*$*-extensions* were introduced: a possible (resp., necessary) i$^*$-extension is a set of arguments that is extension in at least one (resp., every) completion of the iAAF. Similarly, an argument $a$ is accepted under the possible (resp., necessary) perspective if it is accepted in at least one (resp., every) completion.

**Example 1** *Seven agents have been asked for their opinion on which arguments will be claimed in a dispute, and which attacks should be considered. The agents have 7 different views, summarized by the iAAF $IF$ in Figure 1 under the dependency* NAND$(e_1, e_2, e_3)$ *(meaning that, in the agents' opinions, $e_1, e_2, e_3$ do not occur together). So, $IF$ encodes the agent's views as 7 completions, that are the AAFs with the arguments a, b, plus a strict subset of $\{e_1, e_2, e_3\}$, plus the attacks in $IF$ between the arguments in the completions.*

*Under the complete semantics, there is no necessary extension, as no set of arguments is a complete extension in every completion. Examples of possible i*$^*$*-extensions are the sets $S_1 = \{a\}$ (as it is extension in every completion but that where none of $e_1$, $e_2$, $e_3$ occurs) and $S_2 = \{a, b\}$ (as it is extension in the only completion where $e_1$, $e_2$, $e_3$ do not occur). Correspondingly, $a$ is accepted under the possible and the necessary perspective, while $b$ is accepted under the possible perspective, but not under the necessary perspective.*

The main reason for the popularity of iAAFs is their simplicity, mostly deriving from the fact that they do not require a quantitative modeling of the uncertainty. Nevertheless, some

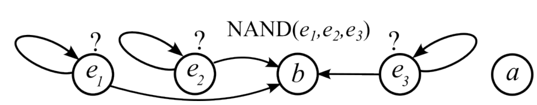


Figure 1: The iAAF of Example 1 ("?" marks uncertain arguments)

quantitative reasoning over iAAFs is likely to be useful when analyzing the dispute modeled by an iAAF. For instance, in Example 1, the fact that $S_1$ and $S_2$ are not necessary i$^*$-extensions but are possible i$^*$-extensions simply tells us that they have some chance to meet the extension's requirements, but $S_1$ and $S_2$ are indistinguishable in terms of how risky it is to consider them as if they were extensions. In this regard, the analysis of the dispute would definitely benefit from knowing the number of completions where $S_1$ and $S_2$ are extensions. In this case, $S_1 = \{a\}$ is an extension in 6 of the 7 completions, while $S_2 = \{a, b\}$ only in 1 of the 7 completions: this tells us that, despite they are both possible i$^*$-extensions, $S_1$ has more chances than $S_2$ to meet the extension's requirement when the actual scenario materializes. Similarly, the information that $b$ is accepted under the possible perspective is much less insightful than knowing that $b$ is accepted in 1 of the 7 completions, which tells us that its acceptance is unlikely. The point is that, in the absence of quantitative measures implying a rank of the completions, it is reasonable to consider them "*alternative scenarios that may occur with the same probability*", and, in turn, to measure the closeness of a set (resp. an argument) to being an extension (resp., accepted) as the percentage of completions where this happens.

In this paper, we present a reasoning paradigm over iAAFs based on three problems, whose input includes an iAAF $IF$, a set of dependencies $\mathcal{D}$, a set of arguments $S$, an argument $a$:
1) PERCVER$^\sigma(IF,\mathcal{D},S)$ and PERCACC$^\sigma(IF,\mathcal{D},a,X)$, asking for the percentages of completions of $IF$ (satisfying $\mathcal{D}$) where $S$ is an extension and $a$ accepted, respectively;
2) CNTCOM$(IF, \mathcal{D})$, counting the completions of $IF$ (satisfying $\mathcal{D}$): this problem supports a preliminary analysis, as its answer is a measure of the uncertainty encoded in the iAAF and helps interpret the answers of PERCVER and PERCACC.

As a first contribution, we study the relationship between our framework and *probabilistic AAFs in the constellations approach* (prAAFs) [Fazzinga *et al.*, 2019; Li *et al.*, 2011; Hunter, 2014; Fazzinga *et al.*, 2022a], that are iAAFs where a probability distribution function is defined over the completions (called *possible worlds* in the context of prAAFs). In particular, we focus on the relationship with the problems PROBVER and PROBACC over prAAFs, that ask for the overall probability of the possible worlds of a prAAF $PF$ where a set is an extension and an argument accepted, respectively. We observe that, although in principle solving PERCVER and PERCACC is the same as solving PROBVER and PROBACC over a suitably constructed prAAF, constructing this "equivalent" prAAF can require a heavy computational cost (besides the fact that the so obtained prAAF is dramatically less compact and less user-friendly than the iAAF). We show that this issue holds also when no dependency is specified, which is somehow counterintuitive, since an iAAF without dependencies seems to be equivalent to a prAAF where arguments and attacks are independent and whose marginal probabilities are set equal to $1/2$ (or some other constant). These results back the need of addressing PERCVER and PERCACC as new distinguished problems. Starting from this, we thoroughly investigate the computational complexity of the quantitative-reasoning paradigm, and show that CNTCOM is #P-complete and PERCVER and PERCACC are FP$^{\#\text{P}}$-complete, with islands of tractability depending on the semantics of extensions and the size of $S$ or the structural properties of the iAAF. The example below helps the reader better appreciate the relevance of studying PERCVER and PERCACC as new problems over iAAFs, distinguished from the classical verification and acceptance problems over prAAFs.

**Example 2** *A group of volunteers have been asked to analyze a text in ancient Greek, whose sentences are the claims of the participants to a philosophical discussion. The text consists of* 10 *sentences. For each sentence, the volunteers have been called for choosing the most accurate among* 3 *alternative translations in English. Moreover, they have been asked to specify the pairs of sentences between which, in their opinion, an attack relationship holds. The result of this process is as follows: for each sentence, each of the* 3 *possible translations has been chosen by some volunteer, and* 92 *different attacks have been specified. In particular,* 90 *attacks have been detected unanimously, while other* 2 *attacks (namely, $\delta_1$, $\delta_2$) have been recognized by some volunteers. As for $\delta_1$ and $\delta_2$, no volunteer has specified that both of them hold.*

*In order to reason on the semantics of the text, the result of the interviews might be represented as a prAAFs: this would require enumerating the subjective views (i.e. the different combinations of the attacks with the alternative translations of the* 10 *sentences), encoding them as distinguished possible worlds, and assigning the same probability to each of them. Indeed, this is a heavy task, since enumerating the possible worlds means considering at least $3^{10}$ combinations of arguments, where some combination must be considered in up to* 3 *variants (that are the ways of combining $\delta_1$ and $\delta_2$ while forbidding their coexistence). On the other hand, resorting to iAAFs as a representation paradigm makes things much easier: the set of alternative possible worlds can be compactly encoded by marking all the arguments providing the translations of the sentences as uncertain, the* 90 *attacks shared by the volunteers as certain, and the two attacks $\delta_1$, $\delta_2$ as uncertain. Then, in order to make the set of completions exactly model the subjective views, it suffices to specify the dependency* NAND$(\delta_1, \delta_2)$ *and, for each sentence, the dependency* CHOICE$(a_1, a_2, a_3)$*, meaning that, in each completion, exactly one of the arguments $a_1, a_2, a_3$ providing alternative translations occurs. Starting from this, without the need of enumerating the subjective views (which would be necessary for both defining the possible worlds and their probability), an analysis of the original text can be accomplished by solving instances of* PERCVER *and* PERCACC*. This way, in the challenging scenario where each sentence has no certain meaning (as witnessed by the existence of different subjective views), the "robustness" of any group of sentences $S$ or any single sentence $s$ can be assessed, in terms of the percentage of subjective views where $S$ is an extension or $s$ is accepted.*

## 2 Preliminaries

**Abstract Argumentation Frameworks (AAFs).** An *abstract argumentation framework* (*AAF*) is a pair $F = \langle A, D\rangle$, where $A$ is a finite set, whose elements are called *arguments*, and $D \subseteq A \times A$ is a binary relation over $A$, whose elements are called *attacks*. Given a set of arguments $S$ and an argument $a$, we say that "$S$ *attacks* $a$" if there is an argument $b$ in $S$ such that $b$ attacks $a$, and that "$a$ *attacks* $S$" if there is an argument $b \in S$ such that $a$ attacks $b$. We say that an argument $a$ (resp., a set of arguments $S$) "*defends* $b$ *against* $c$*'s attack*" if $c$ attacks $b$, while $a$ (resp., $S$) attacks $c$. Moreover, we say that $a$ *is acceptable w.r.t.* $S$ if $S$ defends $a$ against every attack, and that $S$ is *conflict-free* if there is no attack between its arguments. An AAF $F$ can be viewed as a directed graph (called "*argumentation graph*") whose nodes are the arguments and whose edges the attacks.

Several semantics for AAFs have been proposed to identify "reasonable" sets of arguments, called *extensions* [Dung, 1995]. A set $S \subseteq A$ is: an *admissible extension* (ad) iff $S$ is conflict-free and its arguments are acceptable w.r.t. $S$; a *stable extension* (st) iff $S$ is conflict-free and $S$ attacks every $a \in A \setminus S$; a *complete extension* (co) iff $S$ is admissible and contains all the arguments acceptable w.r.t. $S$; a *grounded extension* (gr) iff $S$ is a minimal (w.r.t. $\subseteq$) complete extension; a *preferred extension* (pr) iff $S$ is a maximal (w.r.t. $\subseteq$) complete extension. ad, st, gr, co, pr will be also referred to as *Dungean semantics*. Arguments belonging to at least one (resp., every) extension are said to be *Credulously* (or *Cr*-) accepted (resp., *Skeptically* (or *Sk*-) accepted). The fundamental problems supporting the reasoning over AAFs are the verification that a set $S$ is an extension and that an argument $a$ is $X$-accepted, with $X \in$ {Cr, Sk}.

**Incomplete AAFs (iAAFs)**. iAAFs [Baumeister *et al.*, 2018] allow the uncertainty affecting the real presence of arguments and attacks to be qualitatively modeled.

**Definition 1 (iAAF)** *An* incomplete AAF *(iAAF) is a tuple* $\langle A, A^?, D, D^? \rangle$*, where* $A$ *and* $A^?$ *are disjoint sets of arguments, and* $D$ *and* $D^?$ *disjoint sets of attacks between arguments in* $A \cup A^?$*. The arguments and attacks in* $A$ *and* $D$ *(resp.,* $A^?$ *and* $D^?$*) are said to be* certain *(resp.,* uncertain*), i.e. they are (resp., are not) guaranteed to occur.*

An iAAF compactly represents the alternative scenarios for the argumentation, i.e. all the possible combinations of arguments and attacks that can occur according to what is certain and uncertain. Each scenario is called *completion*.

**Definition 2 (Completion)** *Given an iAAF* $IF = \langle A, A^?, D, D^? \rangle$*, a completion of* $IF$ *is an AAF* $F = \langle A', D' \rangle$ *where* $A \subseteq A' \subseteq (A \cup A^?)$ *and* $D \cap (A' \times A') \subseteq D' \subseteq (D \cup D^?) \cap (A' \times A')$*.*

The possible and necessary perspectives are natural ways to take into account the presence of multiple completions when adapting the notions of *extension* (now renamed "i*-*extension*") and *accepted argument* to iAAFs:

**Definition 3** *Let* $IF$ *be an iAAF,* $S$ *a set of arguments and* $a$ *an argument. Under a semantics* $\sigma$*,* $S$ *is an i*-extension and* $a$ *an* $X$*-accepted argument (with* $X \in$ *{Cr, Sk}) under the possible (resp., necessary) perspective if, for at least one (resp., every) completion* $F$ *of* $IF$*,* $S$ *is an extension of* $F$ *and* $a$ *an* $X$*-accepted argument of* $F$*, respectively.*

In [Fazzinga *et al.*, 2021a; Fazzinga *et al.*, 2021b], the possibility of restricting the set of completions (in order to make it better fit the alternative scenarios) via OR, NAND, CHOICE, IMPLY dependencies involving arguments/attacks was studied (this set was proved to be sufficient to encode any propositional constraint over the completions, expressing which combinations of arguments or of attacks can/cannot occur). Given two sets of arguments $Y$, $Z$, the semantics is:
- $\text{OR}(Y)$: (the completions to be considered are all and only those where) at least one argument in $Y$ occurs,
- $\text{NAND}(Y)$: at least one argument in $Y$ does not occur,
- $\text{CHOICE}(Y)$: exactly one argument in $Y$ occurs,
- $\text{IMPLY}(Y, Z)$: if every argument in $Y$ occurs, then all the arguments in $Z$ occur.

If $Y$ and $Z$ are sets of attacks, the semantics is analogous, but, when checking if a completion satisfies a dependency, only the attacks in $Y$ and $Z$ between arguments that both belong to the completion are considered (this means conditioning the dependencies to the presence of the arguments involved in the attacks). In the presence of a set of dependencies $\mathcal{D}$, the set of completions of an iAAF $IF$ satisfying $\mathcal{D}$ will be denoted as $\mathcal{C}(IF, \mathcal{D})$, and the problems of checking, under the perspective $P \in$ {Possible, Necessary}, if $S$ is an i*-extension and $a$ $X$-accepted (with $X \in$ {Cr, Sk}) as $\text{IVER}^\sigma(IF, \mathcal{D}, S, P)$ and $\text{IACC}^\sigma(IF, \mathcal{D}, a, P, X)$, respectively.

**Probabilistic Abstract Argumentation Frameworks (prAAFs)**. We consider prAAFs following the "*constellations approach*", where probabilities quantitatively model the uncertainty affecting the knowledge of which "possible world" (i.e. combination of the arguments and attacks) actually occurs. In the context of prAAFs, "*possible world*" has the same meaning as "*completion*" in the context of iAAFs, so prAAFs can be viewed as iAAFs where a pdf (*probability distribution function*) is defined over the completions.

**Definition 4** *A prAAF* $PF$ *is a tuple* $\langle A, D, \mathcal{PW}, p \rangle$*, where* $A$ *is a set of arguments,* $D \subseteq A \times A$ *a set of attacks,* $\mathcal{PW}$ *a set of possible worlds over* $A$ *and* $D$*, and* $p$ *is a pdf over* $\mathcal{PW}$*.*

When independence between arguments/attacks is assumed, the pdf $p$ can be encoded compactly, by specifying the marginal probabilities $\mu$ of the arguments/attacks, so that the probability of a possible world $\omega = \langle A', D' \rangle$ is

$$p(\omega) = \Pi_{a \in A'} \mu(a) \times \Pi_{a \in A \setminus A'} \big(1 - \mu(a)\big) \times \Pi_{(a,b) \in D'} \mu\big((a,b)\big) \times \Pi_{(a,b) \in \left(D \cap (A' \times A')\right) \setminus D'} \Big(1 - \mu\big((a,b)\big)\Big), \tag{E1}$$

which means that the marginal probability of an attack $(a, b)$ is interpreted as the probability that the attacks occurs assuming that $a$ and $b$ occur. Thus, under the independence assumption, $\mathcal{PW}$ and $p$ can be replaced by a *marginal probability function* $\mu : A \cup D \to [0, 1]$. In this case, the prAAF will be called IND-prAAF and denoted as a tuple $PF = \langle A, D, \mu \rangle$.

As observed above, prAAFs are a generalization of iAAFs, where the uncertainty is modeled quantitatively, via probabilities. So, when moving from iAAFs to prAAFs, the problems IVER and IACC, that are intrinsically decision problems, become the research problems $\text{PROBVER}^\sigma(PF, S)$ and $\text{PROBACC}^\sigma(PF, a, X)$, asking for the probability that $S$ is an extension and $a$ $X$-accepted (with $X \in$ {Cr, Sk}), i.e. the overall probability of the possible worlds where $S$ is an extension and $a$ $X$-accepted, respectively.

**Counting and functional complexity classes**. We assume familiarity with the complexity classes for decision problems in the polynomial hierarchy PH (in particular, P, NP, coNP, $\Sigma_2^p$ =$\text{NP}^{\text{NP}}$), and recall some complexity classes tailored at counting and functional problems. FP (resp., $\text{FP}^{\mathcal{C}}$) is the class of the functional problems that can be solved in polynomial time by a deterministic Turing machine (resp., by a deterministic Turing machine with an oracle for the class $\mathcal{C}$). #P is the class of the functions $f$ counting the accepting paths of a nondeterministic polynomial-time Turing machine [Valiant, 1979]. More generally, for any class $\mathcal{C}$ in the polynomial hierarchy, #$\mathcal{C}$ is the class of the functions $f$ counting the accepting paths of a nondeterministic polynomial-time with an oracle for the class $\mathcal{C}$ (for instance, #NP is the class of functions counting the accepting paths in a $\text{NP}^{\text{NP}}$ Turing machine). For proving the hardness of counting problems for #P, we resort to *parsimonious reductions*, i.e. (polynomial time) reductions that transform an instance $I'$ of a counting problem to an instance $I''$ of another counting problem such as the answers of $I'$ and $I''$ coincide. As for the hardness for $\text{FP}^{\#\text{P}}$, we use *polynomial time 1-Turing reductions* [Toda and Watanabe, 1992]: a polynomial time 1-Turing reduction from a problem $A$ to a problem $B$ is an algorithm $\mathcal{R}$ that solves $A$ by calling at most once a subroutine solving $B$, such that $\mathcal{R}$ runs in polynomial time if the cost of this subroutine is not considered. It is worth noting that, despite #P and $\text{FP}^{\#\text{P}}$ are different (as the functions in #P return counts, while those in $\text{FP}^{\#\text{P}}$ return more general values), a function is $\text{FP}^{\#\text{P}}$-hard under polynomial time 1-Turing reductions iff it is #P-hard under the same type of reduction. So, sometimes, we prove the hardness for $\text{FP}^{\#\text{P}}$ via reductions from #P-hard problems.

#P- or $\text{FP}^{\#\text{P}}$- completeness mean intractability, since PH$\subseteq$$\text{FP}^{\#\text{P}}$ and the conjecture #P$\not\subseteq$FP is believed to hold (a consequence of its negation is P=NP).

## 3 Quantitative Reasoning Over iAAFs Vs. Probabilistic Reasoning Over prAAFs

We introduce the problems CNTCOM, PERCVER and PERCACC, as the core of a framework for quantitatively reasoning over iAAFs. Then, we investigate their relationship with PROBVER and PROBACC, the probabilistic counterparts of the verification and acceptance problems over prAAFs. In the rest of the paper, an iAAF $IF = \langle A, A^?, D, D^? \rangle$, a set of dependencies $\mathcal{D}$, a set of arguments $S$ of $IF$, and a semantics $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{pr}\}$, are assumed to be given.

**Definition 5** CNTCOM$(IF, \mathcal{D})$ *is the problem of computing* $|\mathcal{C}(IF, \mathcal{D})|$*, and* PERCACC$^\sigma(IF,\mathcal{D},a,X)$ *(with* $X \in \{Cr, Sk\}$*) and* PERCVER$^\sigma(IF,\mathcal{D},S)$ *the problems of computing the percentage of completions in* $\mathcal{C}(IF, \mathcal{D})$ *where, under* $\sigma$*,* $a$ *is* $X$*-accepted and* $S$ *is an extension, respectively. Conventionally, the answer of* PERCVER *and* PERCACC *is* $0$ *if* $\mathcal{C}(IF, \mathcal{D}) = \emptyset$.

Observe that PERCACC$^\sigma(IF,\mathcal{D},a,X)$ trivially returns 0 if $\sigma = \mathtt{ad}$ and $X = Sk$, as no argument can be skeptically accepted in any completion (as $\emptyset$ is always an admissible extension). Hence, in most of the results regarding PERCACC, the combination $\sigma = \mathtt{ad}$ and $X = Sk$ will not be considered.

The following proposition states that PERCVER and PERCACC can be solved by viewing $IF$ as a prAAF whose possible worlds are the completions of $IF$ and whose pdf is the uniform distribution, and then reasoning on this prAAF via PROBVER and PROBACC. Obviously, translating iAAFs to prAAFs makes sense if $\mathcal{D}$ is satisfiable, as prAAFs are not defined if there is no possible world.

**Proposition 1** *Assume* $\mathcal{C}(IF, \mathcal{D}) \neq \emptyset$*. Let* $PF = \langle A \cup A^?, D \cup D^?, \mathcal{PW}, p\rangle$ *be the prAAF where* $\mathcal{PW} = \mathcal{C}(IF, \mathcal{D})$ *and* $p$ *is the pdf that assigns* $1/|\mathcal{PW}|$ *to every possible world of* $PW$*. Then* PERCVER$^\sigma(IF,\mathcal{D},S)$=PROBVER$^\sigma(PF, S)$ *and* PERCACC$^\sigma(IF,\mathcal{D},a,X)$=PROBACC$^\sigma(PF, a, X)$.

Interestingly, although Proposition 1 states that the quantitative reasoning underlying PERCVER and PERCACC can be simulated via classical problems over a prAAF $PF$ having a straightforward logical correspondence with $IF$, this does not diminish the reasonability of addressing PERCVER and PERCACC as new distinguished problems. In fact:

*Issue 1:* prAAFs are not compact, as they require the enumeration of the possible worlds (whose number can be exponential compared with the number of attributes and attacks) and the definition of the pdf over them: this makes their use much less user-friendly than iAAFs, where no pdf must be specified and the completions are not enumerated;

*Issue 2:* prAAFs are more expressive than iAAFs, so a general machinery solving PROBVER/PROBACC may not be capable of exploiting possible simplifications of the implemented computational mechanism allowed by the specificity of the considered prAAFs, that are prAAFs simulating iAAFs;

*Issue 3:* the definition of the pdf $p$ of the prAAF (i.e. the probability value assigned by $p$ to every possible world) is not trivial at all. On the contrary, the theorem below states that CNTCOM$(IF, \mathcal{D})$ (which returns $1/p$) is #P-complete.

**Theorem 1** CNTCOM$(IF, \mathcal{D})$ *is #P-complete, even if* $\mathcal{D} = \emptyset$ *and* $IF$ *contains no certain argument and no certain attack.*

*Proof sketch.* CNTCOM can be viewed as the problem of counting the accepting paths in the non-deterministic polynomial time Turing machine guessing a set of arguments and attacks and checking if it is a completion, so it is in #P. Moreover, it can be shown that it is equivalent to the #P-complete problem of evaluating the overall weight of the homomorphisms between the argumentation graph and a specific weighted graph (see Appendix for details). □

Theorem 1 is of independent interest, since it characterizes a problem that can support the analysis of the real-world modeled by an iAAF: the answer of CNTCOM gives insights into the "extent of uncertainty" encoded in $IF$, and is a baseline that can help better understand the answers of PERCVER and PERCACC. In our current reasoning, Theorem 1 implies that constructing the prAAF $PF$ defined in Proposition 1 is a heavy task (even if the possible worlds were not explicitly enumerated), so it is not reasonable to give up the quantitative framework based on PERCVER and PERCACC with the motivation that it can be simulated by a prAAF with a straightforward logical correspondence with $IF$.

However, one might wonder whether $IF$ can be transformed into an "equivalent" prAAF using a strategy different from ensuring that the possible worlds are equiprobable (that is computationally heavy, as explained above). Obviously, we consider translations towards a prAAF $PF$ that can be easily interpreted by the analyst who is looking into the original iAAF $IF$, so $IF$ should be defined over the same set of arguments as $IF$. In this regard, the following theorem states a strong negative result: under the hypothesis #P$\not\subseteq$ FP, there is no polynomial time transformation of an iAAF into a prAAF over the same arguments that allows for solving any instance of PERCVER via an instance of PROBVER.

**Theorem 2** *Under the hypothesis #P$\not\subseteq$ FP, for every $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{pr}\}$, there is no algorithm taking as input an iAAF $IF$ that runs in polynomial time and translates $IF$ into a prAAF over the same arguments as $IF$ such that, for every set of arguments $S$, $\text{PROBVER}^\sigma(PF, S)$= $\text{PERCVER}^\sigma(IF,\mathcal{D},S)$, or, for each argument $a$, $\text{PERCACC}^\sigma(IF,\mathcal{D},a,X)$= $\text{PROBACC}^\sigma(PF, a, X)$ (except for the case $\sigma = \mathtt{ad}$ and $X = Sk$).*

*(Proof.)* By contradiction, assume that there is an algorithm *trans* whose existence falsifies the part of the statement regarding PERCVER, for some $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{pr}\}$. Let $IF = \langle A, A^?, D, D^? \rangle$ be an iAAF with $A = D = \emptyset$ such that: 1) the completion $C^0$ containing all the arguments in $A^?$ and no attack satisfies $\mathcal{D}$, and 2) $\text{CNTCOM}(IF, \mathcal{D})$ is a hard instance. Such an $IF$ exists, since Theorem 1 guarantees that CNTCOM is #P-hard when $\mathcal{D}$= $\emptyset$, and in this case every completion satisfies the dependencies. Let $PF$ be the IND-prAAF obtained by running *trans* over $IF$. Let $S = A^?$. It is easy to see that the only completion of $IF$ where $S$ is a $\sigma$- extension is $C^0$. Thus, the answer of $\text{PERCVER}^\sigma(IF,\mathcal{D},S)$ is $1/\text{CNTCOM}(IF, \mathcal{D})$. On the other hand, the only possible world of $PF$ where $S$ is a $\sigma$-extension is $\omega_0$, coinciding with $C_0$. Since $\text{PERCVER}^\sigma(IF,\mathcal{D},S)$=$\text{PROBVER}^\sigma(PF, S)$ and $\text{PROBVER}^\sigma(PF, S)$= $p(\omega_0)$ (where $p$ is the pdf of $PF$), we obtain $\text{CNTCOM}(IF, \mathcal{D})$= $1/p(\omega_0)$. Hence the overall procedure (running *trans* and then inverting $p(\omega_0)$) computes $\text{CNTCOM}(IF, \mathcal{D})$ in polynomial time, contradicting (under the assumption #P$\not\subseteq$FP) Theorem 1. The statement for PERCACC can be proved analogously (see Appendix). □

However, the above theorem does not exclude forms of iAAFs for which the polynomial-time translation mentioned in Theorem 2 exists. A nice candidate seems the case where no dependencies are specified: at first sight, one may think that $IF$ can be translated into an IND-prAAF $PF$ over the same arguments and attacks as $IF$, where the marginal probabilities are suitably set so that the possible worlds are equiprobable. Such a translation would also fix *Issue* 1, since the resulting IND-prAAF would have the same size as $IF$ (as now $\mathcal{PW}$ and the pdf $p$ are represented implicitly). Unfortunately, the example below shows that this intuition is wrong.

**Example 3** *The different opinions of 5 agents on the arguments that will occur in a dispute and on the attacks that exist between them are encoded by the 5 completions of the $IF$ having $A = D = \emptyset$, $A^? = \{a, b\}$, and $D^? = \{(a, b)\}$. At first sight, it may seem that reasoning on $IF$ is the same as on the IND-prAAF $PF = \langle A^?, D^?, \mu \rangle$ with $\mu(\cdot) = 1/2$, as this choice for $\mu$ seems to imply that the possible worlds are equiprobable. Indeed, this is false: using Eq.* (E1)*, we obtain that the possible worlds containing both $a$ and $b$ have lower probability than the others. As a matter of fact, there is no $\mu$ that implies a uniform pdf over the possible worlds: the reader can check that the system of equalities whose variables are the arguments' and attacks' marginal probabilities, and whose equations impose that all the possible worlds are assigned the same probability by Eq.* (E1) *has no solution.*

However, also when $\mathcal{D}$= $\emptyset$ (as discussed for the general case above), it is natural to wonder whether there is some general easy-to-compute translation from $IF$ into an equivalent IND-prAAF over the same arguments, that does not pursue the equivalence by making the possible worlds equiprobable. Unfortunately, analogously to the general case, computing such a translation (assuming that it exists) would require a huge computational effort, as formally stated below.

**Theorem 3** *Under the hypothesis #P$\not\subseteq$ FP, even if $\mathcal{D}$= $\emptyset$, under $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{pr}\}$, there is no polynomial-time algorithm translating any iAAF $IF$ into an IND-prAAF $PF$ over the arguments of $IF$ such that, for any set of arguments $S$, $\text{PROBVER}^\sigma(PF, S)$=$\text{PERCVER}^\sigma(IF,\mathcal{D},S)$, or, for any argument $a$, $\text{PERCACC}^\sigma(IF,\mathcal{D},a,X)$= $\text{PROBACC}^\sigma(PF, a, X)$ (except for the case $\sigma = \mathtt{ad}, X = Sk$).*

*(Proof.)* The same strategy as Theorem 2's proof can be used: the difference is that now computing $p(\omega_0)$ requires the polynomial time evaluation of Eq. (E1). □

Our investigation on the translatability of iAAFs in prAAFs continues with the following theorem, describing a form of iAAF for which the polynomial-time translation mentioned in Theorem 3 exists.

**Theorem 4** *Assume that $\mathcal{D}$=$\emptyset$ and every uncertain attack in $IF$ involves at least one certain argument. Let $PF = \langle A \cup A^?, D \cup D^?, \mu \rangle$ be the IND-prAAF where $\forall a \in A\ \mu(a) = 1$, $\forall \delta \in D\ \mu(\delta) = 1$, $\forall \delta \in D^?\ \mu(\delta) = \frac{1}{2}$, and $\forall a \in A^?\ \mu(a) = \frac{2^{|D^?(a)|}}{1+2^{|D^?(a)|}}$, where $D^?(a)$ is the set of uncertain attacks involving $a$. Then, for any $S \subseteq A \cup A^?$, $\text{PERCVER}(IF, \mathcal{D}, S)$ = $\text{PROBVER}^\sigma(PF, S)$, and, for each $a \in A \cup A^?$, $\text{PERCACC}^\sigma(IF,\mathcal{D},a,X)$= $\text{PROBACC}^\sigma(PF, a, X)$.*

*Proof sketch.* It can be proven by showing that $\mu$ implies a uniform pdf over the possible worlds (see Appendix). □

Finally, Theorem 5 below (whose proof is in Appendix) gives an insight into the sufficient condition of Theorem 4, as it states that this condition makes CNTCOM tractable too.

**Theorem 5** $\text{CNTCOM}(IF, \mathcal{D})$ *is in P if $\mathcal{D}$= $\emptyset$ and every uncertain attack of $IF$ involves at least one certain argument.*

**Concluding remarks.** Overall, the results in this section, along with the motivations regarding the usefulness of accomplishing a quantitative reasoning over iAAFs, show the reasonability and relevance of studying PERCVER and PERCACC as new problems: even if there are restricted forms of iAAFs for which there is an easy-to-compute natural representation as a user-friendly prAAF that allows to simulate the reasoning underlying PERCVER and PERCACC, such a translation is not guaranteed to exist in the general case. Starting from this, in the following section we investigate the computational complexity of PERCVER and PERCACC.

# 4 Computational Complexity of Quantitative Reasoning Over iAAFs

We start our investigation by providing a general characterization of PERCVER and PERCACC.

**Theorem 6** PERCACC$^{\sigma}(IF,\mathcal{D},a,X)$ *(except for the case* $\sigma$ = ad *and* $X$ =*Sk) and* PERCVER$^{\sigma}(IF,\mathcal{D},S)$ *are FP*$^{\#P}$*-complete, even if* $\mathcal{D}=\emptyset$.

*Proof.* PERCVER is solved by a polynomial-time Turing machine returning $y/x$, where $x$ is computed by an oracle solving CNTCOM$(IF,\mathcal{D})$, and $y$ by an oracle computing the number of accepting paths in a non-deterministic polynomial-time Turing machine guessing a completion $C$ and then checking if $S$ is a $\sigma$-extension of $C$. The oracle for $x$ is in #P (Theorem 1), while that for $y$ is in #P for $\sigma \in \{\texttt{ad},\texttt{st},\texttt{gr},\texttt{co}\}$ and in #NP for $\sigma=\texttt{pr}$ (since verifying if a set is an extension of a completion is in P and in coNP in the two cases, respectively). This implies that PERCVER is in FP$^{\#P}$ for $\sigma \in \{\texttt{ad},\texttt{st},\texttt{gr},\texttt{co}\}$ and in FP$^{\#NP}$ for $\sigma=\texttt{pr}$. Since, for any $\mathcal{C} \in$ PH, FP$^{\#P}$ coincides (via polynomial-time 1-Turing reductions) with FP$^{\#\mathcal{C}}$ [Toda and Watanabe, 1992], the classes FP$^{\#P}$ and FP$^{\#NP}$ coincide. This implies the membership of PERCVER in FP$^{\#P}$ under every semantics. The membership of PERCACC can be proved similarly, as testing if $a$ is $X$-accepted in a completion is in PH [Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002; Coste-Marquis *et al.*, 2005]. The hardness can be proved via reductions from CNTCOM (see Appendix). □

Starting from this general result, we analyze the sensitivity of the computational complexity to the form of uncertainty encoded in the iAAF and to the semantics of extensions. A first result straightforwardly follows from Theorem 4, which states that, when $\mathcal{D}=\emptyset$ and the uncertain attacks involve at least one certain argument, PERCVER$^{\sigma}(IF,\mathcal{D},S)$ is equivalent to PROBVER$^{\sigma}(PF,S)$, and PERCACC$^{\sigma}(IF,\mathcal{D},a,X)$ to PROBACC$^{\sigma}(PF,a,X)$, where $PF$ is an IND-prAAFs obtained via a polynomial-time transformation from $IF$. Since, under $\sigma \in \{\texttt{ad},\texttt{st}\}$, PROBVER$^{\sigma}(PF,S)$ is in FP over IND-prAAFs (as shown in [Fazzinga *et al.*, 2015]), we obtain:

**Corollary 1** *Assume that* $\mathcal{D}=\emptyset$ *and* $IF$ *is such that every uncertain attack involves at least one certain argument. Under* $\sigma \in \{\texttt{ad},\texttt{st}\}$, PERCVER$^{\sigma}(IF,\mathcal{D},S)$ *is in FP.*

Under the Dungean semantics other than ad and st, PROBVER over IND-prAAFs is not known to be in FP but is FP$^{\#P}$-complete [Fazzinga *et al.*, 2015], and, under every Dungean semantics, PROBACC is FP$^{\#P}$-complete. Thus, Theorem 4 implies that PERCVER is in FP$^{\#P}$ under $\sigma \in \{\texttt{gr},\texttt{co},\texttt{pr}\}$ and PERCACC is in FP$^{\#P}$ under $\sigma \in \{\texttt{ad},\texttt{st},\texttt{gr},\texttt{co},\texttt{pr}\}$. The following theorem states that these are also lower bounds, even if only the arguments or only the attacks are uncertain.

**Theorem 7** *If* $\mathcal{D}=\emptyset$ *and* $IF$ *contains no uncertain argument or no uncertain attacks, then: 1) under* $\sigma \in \{\texttt{gr},\texttt{co},\texttt{pr}\}$, PERCVER$^{\sigma}(IF,\mathcal{D},S)$ *is FP*$^{\#P}$*-complete, and 2) under* $\sigma \in \{\texttt{ad},\texttt{st},\texttt{gr},\texttt{co},\texttt{pr}\}$, PERCACC$^{\sigma}(IF,\mathcal{D},a,X)$ *is FP*$^{\#P}$*-complete (except for the case* $\sigma=$ ad *and* $X=Sk$*).*

*Proof sketch.* The membership is implied by Theorem 6, and the hardness can be shown via reductions from problems counting the truth assignments making positive DNF or CNF formulas true (see Appendix). □

From what shown so far, it is natural to wonder whether the tractability islands of Theorem 5 for CNTCOM and of Corollary 1 for PERCVER can be extended to the case $\mathcal{D} \neq \emptyset$. The following theorem gives a negative answer: even in the simpler case where the uncertainty involves only arguments or only attacks, specifying any form of dependency makes CNTCOM and PERCVER under $\sigma \in \{\texttt{ad},\texttt{st}\}$ intractable.

**Theorem 8** *If* $IF$ *contains no uncertain arguments or no uncertain attacks, and* $\mathcal{D}$ *contains dependencies of only one of the forms* OR, NAND, CHOICE, IMPLY, *then* CNTCOM$(IF,\mathcal{D})$ *is #P-hard and* PERCVER$^{\sigma}(IF,\mathcal{D},S)$ *is FP*$^{\#P}$*-hard.*

Theorem 8 does not mention PERCACC, as it is already FP$^{\#P}$-hard with $\mathcal{D}=\emptyset$ (Theorem 7). As for PERCVER, Theorem 8 interestingly states that PERCVER is FP$^{\#P}$-hard whatever the form of dependencies, even when its decision counterpart (IVER) is tractable (see Section 5 for a comprehensive discussion). We conclude by locating an island of tractability of PERCVER under $\sigma=\texttt{ad}$, that depends on the size of $S$ and how $S$ is connected to the rest of the iAAF. The tractability holds for a form of iAAF for which CNTCOM is not tractable (as no restriction is imposed on the portion of $IF$ outside $S$). To state the result, we denote as *fr*$(S)$ the *frontier* of $S$, i.e. the set of arguments outside $S$ attacking or attacked by $S$. Moreover, we denote as $IF \setminus S$ the iAAF obtained from $IF$ by removing every argument in $S$ and every attack to/from $S$.

**Theorem 9** *Assume that the size of* $S$ *is logarithmic w.r.t. size of* $IF$*, that the arguments in* $\mathrm{fr}(S)$ *are certain, that, for each dependency* $d \in \mathcal{D}$*,* $d$ *involves only arguments/attacks in* $S$ *or only arguments/attacks in* $IF \setminus S$*, and that* $|\mathcal{C}(IF,\mathcal{D})| > 0$*. Then, under* $\sigma=\texttt{ad}$, PERCVER$^{\sigma}(IF,\mathcal{D},S)$ *is in FP.*

*Proof sketch.* As the size of $S$ is logarithmic, the completions of the portion *In* of $IF$ consisting of the arguments and attacks inside $S$ can be enumerated in polynomial time. The rest of the hypothesis avoids the need of evaluating the number the completions of $IF \setminus S$ when computing the percentage of completions of the whole $IF$ where $S$ is an extension, since both the numerator and the denominator of this ratio are proportional to $|\mathcal{C}(IF \setminus S,\mathcal{D})|$, so this term can be simplified, while the other terms can be computed in polynomial time by exploiting the enumeration of the completions of *In*. □

It is worth noting that this tractability result cannot be extended to the other semantics. In fact: 1) under $\sigma \in \{\texttt{co},\texttt{gr},\texttt{pr}\}$, the proof of Theorem 7 (see Appendix) shows that PERCVER is FP$^{\#P}$-hard even in the restricted case where $\mathcal{D}=\emptyset$, *fr*$(S)=\emptyset$, and $S$ is of constant size; 2) under $\sigma=\texttt{st}$, a minor change of the reduction used in the proof of Theorem 6 (consisting in adding no attack involving the fresh argument $a$) shows that PERCVER is FP$^{\#P}$-hard even when $\mathcal{D}=\emptyset$, *fr*$(S)=\emptyset$, and $S$ is of constant size.

The above theorem is the only statement regarding the computational complexity where the satisfiability of the dependencies (i.e. $|\mathcal{C}(IF,\mathcal{D})| > 0$) is a prerequisite. Observe that the general hardness results of Theorem 8 hold even when $\mathcal{D}$ is known to be satisfiable, as they can be proved via parsimonious reductions to CNTCOM from counting problems that always return a value greater than 0. This means that CNTCOM and PERCVER have sources of complexity other than the uncertainty on the existence of at least one completion. In turn, this means that, in the general case, assuming $\mathcal{D}$ satisfiable cannot imply that PERCVER becomes tractable, and makes the result of Theorem 9 relevant.

It is also worth noting that the assumption that $\mathcal{D}$ is satisfiable may be implied in practical scenarios (where the analyst knows that there are some completions satisfying the dependencies) and, in any case, can be preliminarily investigated: in [Fazzinga *et al.*, 2021a], the computational complexity of the satisfiability of dependencies over iAAFs is studied, and shown to be trivial or in P for several combinations of dependencies, although it is NP-complete in the general case.

## 5 Related Work

Except for the case of *Partial Argumentation Frameworks* [Coste-Marquis *et al.*, 2007; Cayrol *et al.*, 2007] and their generalizations [Mailly, 2023], most of the works dealing with incompleteness in AAFs use the completion-based semantics adopted in this paper. In the light of the results in the literature, we can conclude that PERCVER is one of those functional problems that are hard to solve but having an easy decision version (when $\mathcal{D}=\emptyset$). In fact, under all Dungean semantics but $\mathtt{pr}$, PERCVER is $FP^{\#P}$-hard, while, as shown in [Fazzinga *et al.*, 2020], deciding the existence of a completion where $S$ is an extension is in P (under $\sigma = \mathtt{pr}$, the decision problem is instead $\Sigma_2^p$-complete). On the other hand, in [Baumeister *et al.*, 2021], the acceptance problem was shown to be hard already in the decision version, that ranges from NP-complete to $\Sigma_3^p$-complete, depending on the semantics and the perspective (*Cr* or *Sk*), except for the trivial case of the skeptical acceptance under $\sigma = \mathtt{ad}$. As for the case $\mathcal{D}\neq\emptyset$, in [Fazzinga *et al.*, 2021b; Fazzinga *et al.*, 2021a; Mailly, 2021] it was shown that specifying dependencies (OR, NAND, CHOICE, IMPLY) typically makes the complexity of the verification increase, for the semantics under which it was polynomial, to NP-complete. There are notable exceptions: the verification remains in P under $\sigma \in \{\mathtt{ad}, \mathtt{st}\}$ when only arguments or only attacks are uncertain, and only OR or combinations of NAND and IMPLY are used. In this regard, we have shown that $\text{PERCVER}^{\sigma}(IF,\mathcal{D},S)$ is $FP^{\#P}$-complete under every semantics even if only arguments or attacks are uncertain and dependencies of only one form are used.

As for the relationship, in terms of complexity, with prAAFs in the constellations approach [Li *et al.*, 2011], comparing our results with those on PROBVER and PROBACC in [Fazzinga *et al.*, 2019; Fazzinga *et al.*, 2015], we can draw these conclusions:

- if $\mathcal{D} = \emptyset$ (when it makes sense to compare iAAFs with IND-prAAFs), some aspects of the quantitative reasoning over iAAFs are more complex than over IND-prAAFs: under $\sigma \in \{\mathtt{ad}, \mathtt{st}\}$, PERCVER is $FP^{\#P}$-complete while PROBVER is in FP. Under the other semantics, the complexity of PERCVER and PROBVER is the same. Instead, PERCACC and PROBACC have the same complexity under every Dungean semantics.

– if $\mathcal{D} \neq \emptyset$ (when it makes sense to compare iAAFs with prAAFs without the independence assumption), both PERCVER and PERCACC are $FP^{\#P}$-complete, so the complexity is higher than PROBVER (that is in FP under $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{gr}, \mathtt{co}\}$ and $FP^{||NP}$-complete under $\sigma = \mathtt{pr}$) and than PROBACC (that, under each semantics $\sigma$, is in $FP^{||\mathcal{C}_\sigma}$, where $\mathcal{C}_\sigma$ is the complexity class of the classical acceptance problem ACC under $\sigma$). Indeed, this was quite expected, since in prAAFs the possible worlds are enumerated, so the computational complexity benefits from some "discount" compared with iAAFs, where the completions are compactly encoded.

This work is also related to [Alfano *et al.*, 2023a], where some relationships between prAAFs and iAAFs have been analyzed, and with prAAFs following the *epistemic approach* [Thimm, 2012; Hunter and Thimm, 2014], where, similarly to IND-prAAFs, arguments are associated with probabilities. However, these probabilities are degrees of belief in the acceptance of the arguments, so the reasoning paradigm does not take into account the possibility that the structure of the argumentation graph changes as the effect of considering the presence/absence of arguments or attacks.

Our framework is also related to the use of constraints, preferences, and/or explicit acceptance conditions to support the reasoning (as done in several frameworks where these constructs are used to filter out extensions [Alfano *et al.*, 2023a; Alfano *et al.*, 2023b; Alfano *et al.*, 2023d; Alfano *et al.*, 2023c; Alfano *et al.*, 2024], rather than completions), to iAAFs with supports [Fazzinga *et al.*, 2018; Fazzinga *et al.*, 2023], to the credulous/skeptical conclusion problems in *Control Argumentation Frameworks* [Dimopoulos *et al.*, 2018], and to revising AAFs to enforce the existence of an extension [Baumann and Ulbricht, 2019], or to make a set an extension [Coste-Marquis *et al.*, 2015]. In this regard, if we interpret the uncertain arguments/attacks as elements that can be added or removed, the framework can be used to count the possible ways of making a set an extension (or an argument accepted) via insertions or removals.

Finally, this work is related to [Baroni *et al.*, 2010; Dewoprabowo *et al.*, 2022; Fichte *et al.*, 2019], where the problem of counting the number of extensions in an AAF has been addressed. A nice direction of future work is integrating these paradigms with our framework into a new reasoning paradigm addressing questions like "*How many sets of arguments are extensions in at least 80% of the completions?*"

## 6 Conclusions and Future Work

A new quantitative reasoning paradigm supporting the analysis of iAAFs has been introduced. Its core consists in the problems PERCVER and PERCACC, whose answers give insights into how likely it is that, once the dispute modeled by the iAAF takes place, a set is an extension or an argument is accepted. The relationship with the classical reasoning paradigm over prAAFs has been investigated, first in terms of translatability, and then of computational complexity, after having characterized this aspect for the introduced framework. Future work will investigate the possibility of simultaneously looking into the counts of completions and extensions, as sketched at the end of the previous section.

## Acknowledgements

We acknowledge partial financial support from MUR project PRIN 2022 EPICA (CUP H53D23003660006) funded by the European Union - Next Generation EU and from the PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the NextGenerationEU.

## References

[Alfano *et al.*, 2023a] Gianvincenzo Alfano, Marco Calautti, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Explainable acceptance in probabilistic and incomplete abstract argumentation frameworks. *Artif. Intell.*, 323:103967, 2023.

[Alfano *et al.*, 2023b] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Abstract argumentation framework with conditional preferences. In *Proc. of the AAAI Conf.*, pages 6218–6227, 2023.

[Alfano *et al.*, 2023c] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. On acceptance conditions in abstract argumentation frameworks. *Inf. Sci.*, 625:757–779, 2023.

[Alfano *et al.*, 2023d] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Preferences and constraints in abstract argumentation. In *Proc. of the IJCAI Conf.*, pages 3095–3103, 2023.

[Alfano *et al.*, 2024] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Complexity of credulous and skeptical acceptance in epistemic argumentation framework. In *Proc. of the AAAI Conf.*, pages 10423–10432, 2024.

[Baroni *et al.*, 2010] Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On extension counting problems in argumentation frameworks. In Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo Ricardo Simari, editors, *Proc. Computational Models of Argument (COMMA), Desenzano del Garda, Italy*, pages 63–74, 2010.

[Baumann and Ulbricht, 2019] Ringo Baumann and Markus Ulbricht. If nothing is accepted - repairing argumentation frameworks. *J. Artif. Intell. Res.*, 66:1099–1145, 2019.

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

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

[Cayrol *et al.*, 2007] Claudette Cayrol, Caroline Devred, and Marie-Christine Lagasquie-Schiex. Handling ignorance in argumentation: Semantics of partial argumentation frameworks. In *Proc. Conf. Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU), Hammamet, Tunisia*, pages 259–270, 2007.

[Coste-Marquis *et al.*, 2005] Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. Symmetric argumentation frameworks. In *Proc. of ECSQARU*, pages 317–328, 2005.

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

[Coste-Marquis *et al.*, 2015] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. Extension enforcement in abstract argumentation as an optimization problem. In *Proc. Int. Conf. on Artificial Intelligence (IJCAI)*, pages 2876—-2882, 2015.

[Dewoprabowo *et al.*, 2022] Ridhwan Dewoprabowo, Johannes Klaus Fichte, Piotr Jerzy Gorczyca, and Markus Hecher. A practical account into counting dung's extensions by dynamic programming. In *Proc. Logic Programming and Nonmonotonic Reasoning (LPNMR), Genova, Italy*, pages 387–400, 2022.

[Dimopoulos and Torres, 1996] Yannis Dimopoulos and Alberto Torres. Graph theoretical structures in logic programs and default theories. *Theor. Comput. Sci.*, 170(1-2):209–244, 1996.

[Dimopoulos *et al.*, 2018] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Control argumentation frameworks. In *Proc. Conf. on Artificial Intelligence (AAAI), New Orleans, USA*, pages 4678–4685, 2018.

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

[Dunne and Bench-Capon, 2002] Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artif. Intell.*, 141(1/2):187–203, 2002.

[Fazzinga *et al.*, 2015] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log. (TOCL)*, 16(3):22, 2015.

[Fazzinga *et al.*, 2018] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Probabilistic bipolar abstract argumentation frameworks: complexity results. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1803–1809, 2018.

[Fazzinga *et al.*, 2019] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence. *Artif. Intell.*, 268:1–29, 2019.

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

[Fazzinga *et al.*, 2021a] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over argument-incomplete aafs in the presence of correlations. In Zhi-Hua Zhou, editor, *Proc. International Joint Conference on Artificial Intelligence (IJCAI), Virtual Event / Montreal, Canada*, pages 189–195, 2021.

[Fazzinga *et al.*, 2021b] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over attack-incomplete aafs in the presence of correlations. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proc. International Conference on Principles of Knowledge Representation and Reasoning (KR), Online event*, pages 301–311, 2021.

[Fazzinga *et al.*, 2022a] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Abstract argumentation frameworks with marginal probabilities. In *Proc. Int. Joint Conference on Artificial Intelligence (IJCAI), Vienna, Austria*, pages 2613–2619, 2022.

[Fazzinga *et al.*, 2022b] Bettina Fazzinga, Sergio Flesca, Filippo Furfaro, and Luigi Pontieri. Process mining meets argumentation: Explainable interpretations of low-level event logs via abstract argumentation. *Inf. Syst.*, 107:101987, 2022.

[Fazzinga *et al.*, 2022c] Bettina Fazzinga, Andrea Galassi, and Paolo Torroni. A privacy-preserving dialogue system based on argumentation. *Intell. Syst. Appl.*, 16:200113, 2022.

[Fazzinga *et al.*, 2023] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Incomplete bipolar argumentation frameworks. In *Proc. European Conference on Artificial Intelligence (ECAI), Kraków, Poland*, pages 684–691, 2023.

[Fichte *et al.*, 2019] Johannes Klaus Fichte, Markus Hecher, and Arne Meier. Counting complexity for reasoning in abstract argumentation. In *Proc. Conference on Artificial Intelligence (AAAI), Honolulu, Hawaii, USA*, pages 2827–2834, 2019.

[Hunter and Thimm, 2014] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with epistemic extensions. In *Proc. of DARe@ECAI*, 2014.

[Hunter, 2014] Anthony Hunter. Probabilistic qualification of attack in abstract argumentation. *Int. J. Approx. Reasoning*, 55(2):607–638, 2014.

[Li *et al.*, 2011] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proc.of TAFA*, 2011.

[Mailly, 2021] Jean-Guy Mailly. Constrained incomplete argumentation frameworks. In Jirina Vejnarová and Nic Wilson, editors, *Proc. Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQUARU), Prague, Czech Republic*, pages 103–116, 2021.

[Mailly, 2023] Jean-Guy Mailly. Extension-based semantics for incomplete argumentation frameworks: properties, complexity and algorithms. *J. Log. Comput.*, 33(2):406–435, 2023.

[Thimm, 2012] Matthias Thimm. A probabilistic semantics for abstract argumentation. In *Proc. of ECAI*, pages 750–755, 2012.

[Toda and Watanabe, 1992] Seinosuke Toda and Osamu Watanabe. Polynomial Time 1-Turing Reductions from #PH to #P. *Theor. Comput. Sci.*, 100(1):205–221, 1992.

[Valiant, 1979] Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci. (TCS)*, 8:189–201, 1979.