Credulous acceptability in probabilistic abstract argumentation: complexity results

Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro

DIMES, University of Calabria, Italy
ICAR-CNR, Italy
{flesca,furfaro}@dimes.unical.it
{fazzinga}@icar.cnr.it

Abstract. Probabilistic abstract argumentation combines Dung's abstract argumentation framework with probability theory in order to model uncertainty in argumentation. In this setting, we address the fundamental problem of computing the probability that an argument is (credulously) acceptable according to a given semantics. Specifically, we focus on the most popular semantics (i.e., *admissible*, *stable*, *semi-stable*, *complete*, *grounded*, *preferred*, *ideal*, *ideal-set*), and show that computing the probability that an argument is credulously accepted is $FP^{\#P}$-complete independently from the adopted semantics.

Keywords: Acceptability, Probabilistic Abstract Argumentation, Complexity

1 Introduction

Several argumentation frameworks have been proposed, with the aim of suitably modeling disputes between two or more parties. Typically, argumentation frameworks model both the possibility of parties to present arguments supporting their theses, and the possibility that some arguments rebut other arguments.

A powerful yet simple argumentation framework is that proposed in the seminal paper [10], called *abstract argumentation framework* (AAF). An AAF is a representation of a dispute in terms of an *argumentation graph* $\langle A, D\rangle$, where A is the set of nodes (each called *argument*) and D is the set of edges (each called *defeat* or, equivalently, *attack*). Basically, an argument is an abstract entity that may attack and/or be attacked by other arguments, and an attack expresses the fact that an argument rebuts/weakens another argument.

Several semantics for AAFs, such as *admissible*, *stable*, *preferred*, *complete*, *grounded*, *semi-stable*, *ideal-set* and *ideal*, have been proposed [10, 11, 2, 4] to identify "reasonable" sets of arguments, called *extensions*. Basically, each of these semantics corresponds to some properties that "certify" whether a set of arguments can be profitably used to support a point of view in a discussion. For instance, under the *admissible* semantics, a set S of arguments is an extension if S is "*conflict-free*" (i.e., there is no defeat between arguments in S) and is "robust" against the other arguments (i.e., every argument outside S attacking an argument in S is counterattacked by an argument in S). This means that one using the set of arguments S in a discussion does not contradict her/himself, and can rebut to the arguments possibly presented by the other parties.

As a matter of fact, in the real world, arguments and defeats are often *uncertain*. For instance, consider an argument a (or a defeat δ) encoding an interpretation or a translation of the description of a fact reported in a reference text. Then, a (or δ) may be uncertain in the sense that the original paragraph may have interpretations other than that encoded by a (or δ).

Thus, several proposals have been made to model uncertainty in AAFs, by considering weights, preferences, or probabilities associated with arguments and/or defeats. One of the most popular approaches based on probability theory for modeling the uncertainty is the so called *constellations approach* [12, 32, 7, 9, 25, 27, 30, 19, 20]: the dispute is represented by means of a *Probabilistic Argumentation Framework* (PrAAF), that consists in a set of alternative scenarios, each represented by an argumentation graph associated with a probability. The various works in the literature investigating PrAAFs can differ in the assumption on how the probability distribution function (pdf) over the scenarios is specified. For instance, in [12], the pdf is defined "extensively", by enumerating all the possible scenarios and, for each of them, the value of its probability. On the other hand, in [30], the restriction that arguments and defeats are independent is assumed, and this is exploited to simplify the way probabilities are specified: the pdf is not explicitly specified, as it is implied by the marginal probabilities associated with arguments and defeats.

As regards reasoning over an argumentation framework, in the deterministic setting, two classical problems supporting the reasoning over AAFs have been deeply investigated and used in practical applications:

- $\text{EXT}^{sem}(S)$: the problem of deciding whether a set of arguments S is an extension according to the semantics *sem*;
- CA^{sem}: the problem of deciding whether the argument a is *acceptable*, i.e., it belongs to at least one extension under the semantics *sem*.

Basically, the relevance of these problems is that solving $\text{EXT}^{sem}(S)$ supports the decision on whether presenting a set of arguments in a dispute is a reasonable strategy, while solving CA^{sem} focuses this analysis on single arguments. In the probabilistic setting, there are multiple scenarios to be taken into account, and an argument a can be acceptable in some scenarios, but not in others. Thus, the natural probabilistic counterparts P-$\text{EXT}^{sem}(S)$ and P-CA^{sem} of the above-mentioned problems $\text{EXT}^{sem}(S)$ and CA^{sem} consist in evaluating the overall probability that S is an extension and a is acceptable, respectively, where "overall" means summing the probabilities of the scenarios where the property is verified.

The complexity of $\text{EXT}^{sem}(S)$ and P-$\text{EXT}^{sem}(S)$ has been deeply investigated. Specifically, in the probabilistic setting the complexity of P-$\text{EXT}^{sem}(S)$ has been investigated both in the case that the assumption that arguments and defeats are independent holds [19–22] and in more general cases [23].

As regards CA^{sem}, its complexity has been deeply investigated, showing that it ranges from $PTIME$ to $\sum_2^p$-complete, depending on the adopted semantics [6, 14, 5, 13]. Detailed complexity results from the literature are reported in the first column of Table 1.

However, to the best of our knowledge no work has been done for characterizing the complexity of P-CA^{sem}.

In this paper we focus on the constellations approach with independence proposed in [30] and analyze the complexity of P-CA^{sem} showing that it is $FP^{\#P}$-complete independently from the adopted semantics (see the second column of Table 1). Proving that the problem is intractable backs the usage of alternative strategies for solving P-CA^{sem}, such as resorting to Monte-carlo based probability estimation approaches.

Table 1: Complexity of CA^{sem}and P-CA^{sem}.

sem	CA^{sem}	P-CA^{sem}
admissible	NP-complete	$FP^{\#P}$-complete
stable	NP-complete	$FP^{\#P}$-complete
complete	NP-complete	$FP^{\#P}$-complete
grounded	$PTIME$	$FP^{\#P}$-complete
semi-stable	$\sum_2^p$-complete	$FP^{\#P}$-complete
preferred	NP-complete	$FP^{\#P}$-complete
ideal-set	in Θ_2^p, $coNP$-h	$FP^{\#P}$-complete
ideal	in Θ_2^p, $coNP$-h	$FP^{\#P}$-complete

2 Preliminaries

In this section, we briefly recall some basic notions about computational complexity, and concisely overview Dung's abstract argumentation framework, and its probabilistic extension introduced in [30].

2.1 Complexity

The computational complexity of the problem addressed in this paper is related to the complexity classes of counting problems. A counting problem f is a function from strings over a finite alphabet into integers. $\#P$ is the complexity class of the functions f such that f counts the number of accepting paths of a nondeterministic polynomial-time Turing machine [34]. Although the problem addressed in the paper is closely related to $\#P$, strictly speaking, it cannot belong to it, since the outputs of our problem are not integers. In fact, we deal with the class $FP^{\#P}$, that is, the class of functions computable by a polynomial-time Turing machine with a $\#P$ oracle. In this regard, note that a function is

$FP^{\#P}$-hard iff it is $\#P$-hard, and thus to prove that a problem is $FP^{\#P}$-hard it suffices to reduce a $\#P$-hard problem to it.

2.2 Abstract Argumentation

An *abstract argumentation framework* [10] (AAF) is a pair $\langle A, D\rangle$, where A is a finite set, whose elements are referred to as *arguments*, and $D \subseteq A \times A$ is a binary relation over A, whose elements are referred to as *defeats* (or *attacks*). An argument is an abstract entity whose role is entirely determined by its relationships with other arguments. Given an AAF $\mathcal{A}$, we also refer to the set of its arguments and the set of its defeats as $Arg(\mathcal{A})$ and $Def(\mathcal{A})$, respectively.

Given arguments $a, b \in A$, we say that a *defeats* b iff there is $(a, b) \in D$. Similarly, a set $S \subseteq A$ *defeats* an argument $b \in A$ iff there is $a \in S$ such that a *defeats* b.

A set $S \subseteq A$ of arguments is said to be *conflict-free* if there are no $a, b \in S$ such that a *defeats* b. An argument a is said to be *acceptable* w.r.t. $S \subseteq A$ iff $\forall b \in A$ such that b *defeats* a, there is $c \in S$ such that c *defeats* b. Given a set $S \subseteq A$ of arguments, we define S^+ as the set of arguments that are defeated by S, that is, $S^+ = \{a \in A \text{ s.t. } S \text{ defeats } a\}$.

Several semantics for AAFs have been proposed to identify "reasonable" sets of arguments, called *extensions*. We consider the following well-known semantics [10, 11, 4]: *admissible* (`ad`), *stable* (`st`), *complete* (`co`), *grounded* (`gr`), *semi-stable* (`sst`), *preferred* (`pr`), *ideal-set* (`ids`), and *ideal* (`ide`).
A set $S \subseteq A$ is said to be:

- an *admissible extension* iff S is conflict-free and all its arguments are acceptable w.r.t. S;
- a *stable extension* iff S is conflict-free and S defeats each argument in $A \setminus S$;
- a *complete extension* iff S is admissible and S contains all the arguments that are acceptable w.r.t. S;
- a *grounded extension* iff S is a minimal (w.r.t. $\subseteq$) complete set of arguments;
- a *semi-stable extension* iff S is a complete extension where $S \cup S^+$ is maximal (w.r.t. $\subseteq$);
- a *preferred extension* iff S is a maximal (w.r.t. $\subseteq$) admissible set of arguments;
- an *ideal-set extension* iff S is admissible and S is contained in every preferred set of arguments;
- an *ideal extension* iff S is a maximal (w.r.t. $\subseteq$) ideal-set extension;

Example 1. Consider the AAF $\langle A,\ D\rangle$, where the set A of arguments is $\{a, b, c\}$ and the set D of defeats is $\{\delta_1 = (b, a), \delta_2 = (b, c), \delta_3 = (c, b)\}$, where δ_2 and δ_3 encode the fact that arguments b and c attack each other. A graphical rapresentation of the AAF $\langle A,\ D\rangle$ is reported in Figure 1, where arguments are represented as nodes and defeats are represented as edges between nodes. As $S = \{a, c\}$ is conflict-free and every argument in S is acceptable w.r.t. S, it is the case that S is admissible. It is easy to see that the sets $\{b\}$, $\{c\}$, and $\emptyset$

are admissible extensions as well. Since S is conflict-free and defeats b (the only argument in $A\backslash S$) it is also a stable extension. As S is maximally admissible, it a preferred extension, while $\{c\}$ is not, since a is acceptable w.r.t $\{c\}$. It is easy to check that S is complete, while it neither is grounded (since it is not minimally complete, as the emptyset is complete) nor ideal or ideal-set (since it is not contained in $\{b\}$ which is a preferred extension). $\square$

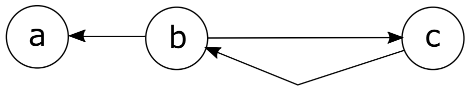

Fig. 1: The AAF $\langle A,\, D\rangle$

Given an AAF $\mathcal{A}$, a set $S \subseteq Arg(\mathcal{A})$ of arguments, an argument $a \in Arg(\mathcal{A})$ and a semantics sem $\in$$\{$`ad`, `st`, `gr`, `co`, `sst`, `pr`, `ids`, `ide`$\}$, we define the function $acc(\mathcal{A}, sem, a)$ that returns `true` if $\exists S \subseteq Arg(\mathcal{A})$ such that $a \in S$ and S is an extension according to sem, `false` otherwise.

Example 2. Consider the AAF $\langle A,\, D\rangle$ introduced in Example 1. The argument a is credulously accepted under the admissible, stable, complete, semi-stable and preferred semantics. a is not credulousy accepted under the grounded, ideal-set and ideal semantics. $\square$

2.3 Probabilistic Abstract Argumentation

We now review the *probabilistic* abstract argumentation framework proposed in [30].

Definition 1 (PrAAF). *A probabilistic argumentation framework (PrAAF) is a tuple $\langle A, P_A, D, P_D\rangle$ where $\langle A, D\rangle$ is an AAF, and P_A and P_D are, respectively, functions assigning a non-zero*[1] *probability value to each argument in A and defeat in D, that is, $P_A : A \to (0, 1] \cap \mathbb{Q}$ and $P_D : D \to (0, 1] \cap \mathbb{Q}$.*

Basically, the value assigned by P_A to an argument a represents the probability that a actually occurs, whereas the value assigned by P_D to a defeat (a, b) represents the conditional probability that a defeats b given that both a and b occur.

The meaning of a PrAAF is given in terms of possible worlds, each of them representing a scenario that may occur in the reality. Given a PrAAF $\mathcal{F}$, a possible world is modeled by an AAF which is derived from $\mathcal{F}$ by considering only a subset of its arguments and defeats. More formally, given a PrAAF $\mathcal{F} = \langle A, P_A, D, P_D\rangle$, a possible world w of $\mathcal{F}$ is an AAF $\langle A', D'\rangle$ such that $A' \subseteq A$ and $D' \subseteq D \cap (A' \times A')$. The set of the possible worlds of $\mathcal{F}$ will be denoted as $pw(\mathcal{F})$.

[1] Assigning probability equal to 0 to arguments/defeats is useless.

Example 3. As a running example, consider the PrAAF $\mathcal{F} = \langle A, P_A, D, P_D \rangle$ reported in Figure 2, where A and D are those of Example 1, and $P_A(a) = .9, P_A(b) = .7, P_A(c) = .2, P_D(\delta_1) = .9$ and $P_D(\delta_2) = P_D(\delta_3) = 1$. The set $pw(\mathcal{F})$ contains the following possible worlds:

$w_1 = \langle \emptyset, \emptyset \rangle$	$w_2 = \langle \{a\}, \emptyset \rangle$	$w_3 = \langle \{b\}, \emptyset \rangle$	$w_4 = \langle \{c\}, \emptyset \rangle$
$w_5 = \langle \{a, b\}, \emptyset \rangle$	$w_6 = \langle \{a, c\}, \emptyset \rangle$	$w_7 = \langle \{b, c\}, \emptyset \rangle$	$w_8 = \langle A, \emptyset \rangle$
$w_9 = \langle \{a, b\}, \{\delta_1\} \rangle$	$w_{10} = \langle \{b, c\}, \{\delta_3\} \rangle$	$w_{11} = \langle \{b, c\}, \{\delta_2\} \rangle$	$w_{12} = \langle \{b, c\}, \{\delta_2, \delta_3\} \rangle$
$w_{13} = \langle A, \{\delta_1\} \rangle$	$w_{14} = \langle A, \{\delta_1, \delta_3\} \rangle$	$w_{15} = \langle A, \{\delta_1, \delta_2\} \rangle$	$w_{16} = \langle A, D \rangle$
$w_{17} = \langle A, \{\delta_2\} \rangle$	$w_{18} = \langle A, \{\delta_3\} \rangle$	$w_{19} = \langle A, \{\delta_2, \delta_3\} \rangle$	

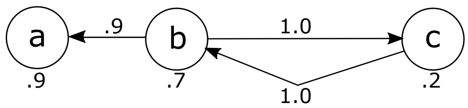

Fig. 2: The PrAAF $\langle A, P_A, D, P_D \rangle$

□

An interpretation for a PrAAF $\mathcal{F} = \langle A, P_A, D, P_D \rangle$ is a probability distribution function I over the set $pw(\mathcal{F})$ of the possible worlds. Assuming that arguments represent pairwise independent events, and that each defeat represents an event conditioned by the occurrence of its argument events but independent from any other event, the interpretation for the PrAAF $\mathcal{F} = \langle A, P_A, D, P_D \rangle$ is as follows. For each possible world $w \in pw(\mathcal{F})$, w is assigned by I the probability:

$$
\begin{aligned}
I(w) = & \prod_{a \in Arg(w)} P_A(a) \times \prod_{a \in A \setminus Arg(w)} (1 - P_A(a)) \times \\
& \times \prod_{\delta \in Def(w)} P_D(\delta) \times \prod_{\delta \in \overline{D}(w) \setminus Def(w)} (1 - P_D(\delta))
\end{aligned}
$$

where $\overline{D}(w)$ is the set of defeats that may appear in the possible world w, that is $\overline{D}(w) = D \cap (Arg(w) \times Arg(w))$. Hence, the probability of a possible world w is given by the product of four contributions: (i) the product of the probabilities of the arguments belonging to w; (ii) the product of the one's complements of the probabilities of the arguments that do not appear in w; (iii) the product of the conditional probabilities of the defeats in w (recall that a defeat $\delta = (a, b)$ may appear in w only if both a and b are in w); (iv) the product of the one's complements of the conditional probabilities of the defeats that, though they may appear in w, they do not.

Example 4. Continuing our running example, the interpretation I for $\mathcal{F}$ is as follows. First of all, observe that, for each possible world $w \in pw(\mathcal{F})$, if both arguments b and c belong to $Arg(w)$ and $\delta_2 \notin Def(w)$ or $\delta_3 \notin Def(w)$, then $I(w) = 0$. The probabilities of the other possible worlds are the following:

$I(w_1) = .024$	$I(w_2) = .216$	$I(w_3) = .056$	$I(w_4) = .006$
$I(w_5) = .0504$	$I(w_6) = .054$	$I(w_9) = .4536$	$I(w_{12}) = .014$
$I(w_{16}) = .1134$	$I(w_{19}) = .0126$		

□

The probability that an argument a is acceptable according to a given semantics *sem* is defined as the sum of the probabilities of the possible worlds w for which a is acceptable according to *sem*, i.e., $acc(w, sem, a) = \texttt{true}$.

Definition 2 **($PrCA^{sem}_{\mathcal{F}}(a)$).** *Given a PrAAF $\mathcal{F} = \langle A, P_A, D, P_D \rangle$, an argument $a \in A$, and a semantics sem, the probability $PrCA^{sem}_{\mathcal{F}}(a)$ that a is acceptable under sem is*

$$PrCA^{sem}_{\mathcal{F}}(a) = \sum_{w \in pwAcc(\mathcal{F},a)} I(w),$$

where $pwAcc(\mathcal{F}, a) = \{w \in pw(\mathcal{F}) \,|\, acc(w, sem, a) = \textit{true}\}$.

The following example shows usages of this definition.

Example 5. In our running example, the probabilities that the arguments a and b are acceptable w.r.t. the admissible semantics are as follows:

$$PrCA^{\mathtt{ad}}_{\mathcal{F}}(a) = I(w_2) + I(w_5) + I(w_6) + I(w_{16}) + I(w_{19}) = .4464$$

$$PrCA^{\mathtt{ad}}_{\mathcal{F}}(b) = I(w_3) + I(w_5) + I(w_9) + I(w_{12}) + I(w_{16}) + I(w_{19}) = .8134$$

□

Obviously, computing $PrCA^{sem}_{\mathcal{F}}(a)$ by directly applying Definition 2 would require exponential time, since it relies on summing the probabilities of an exponential number of possible worlds. However, this does not rule out the possibility that efficient strategies for computing $PrCA^{sem}_{\mathcal{F}}(a)$ exist. Unfortunately, this is not true in the general case. Indeed, in the next section we show that the problem of computing $PrCA^{sem}_{\mathcal{F}}(S)$ is $FP^{\#P}$-complete.

Definition 3 **(P-CA^{sem}).** *Given a PrAAF $\mathcal{F} = \langle A, P_A, D, P_D \rangle$, an argument $a \in A$, and a semantics sem,* P-CA^{sem} *is the problem of computing the probability $PrCA^{sem}_{\mathcal{F}}(a)$.*

3 Complexity of credulous acceptability in probabilistic abstract argumentation

In this section we characterize the complexity of P-CA^{sem}, by first showing that it is in $FP^{\#P}$ and then showing that it is $FP^{\#P}$-hard.

3.1 Upper bound

Theorem 1. *For sem* $\in$ {`ad`, `st`, `co`, `gr`, `sst`, `pr`, `ids`, `ide` }, *it holds that* P-CA^{sem} *is in* $FP^{\#P}$.

Proof. First, observe that $PrCA_{\mathcal{F}}^{sem}(a)$ with *sem* $\in$ {`ad`, `st`, `co`, `gr`, `sst`, `pr`, `ids`, `ide` }, can be expressed as a rational number whose denominator d is the product of the denominators of the probabilities of arguments in A and defeats in D.

We first prove the statement for *sem* = `gr`. First observe that CA^{sem} is in *PTIME* for *sem* = `gr`. We show that $\mathrm{PrCA}_{\mathcal{F}}^{\texttt{gr}}(a)$ can be computed by a polynomial time algorithm $\mathcal{A}$ with access to a $\#P$ oracle. Algorithm $\mathcal{A}$ first computes d in polynomial time w.r.t. the size of $\mathcal{F}$, then calls a $\#P$ oracle to determine the numerator of $\mathrm{PrCA}_{\mathcal{F}}^{\texttt{gr}}(a)$. The oracle counts the number of accepting paths of a nondeterministic polynomial-time Turing machine M such that:

(i) M nondeterministically guesses a subset of arguments in A and defeats in D so that each leaf of the resulting computation tree is a possible world $w \in pw(\mathcal{F})$;
(ii) At each leaf, let w be the guessed world, and $I(w)$ its probability, the computation tree is then split again $d \cdot I(w)$ times to reflect the probability of the guessed world (for each $w \in pw(\mathcal{F})$, $I(w)$ is a rational number whose denominator is d, and $I(w)$ can be computed in polynomial time w.r.t. the size of $\mathcal{F}$);
(iii) Finally, M checks in polynomial time if a is an acceptable argument in the world w w.r.t. the grouded semantics.

Thus, the number of accepting paths of M is $d\cdot$ $\mathrm{PrCA}_{\mathcal{F}}^{\texttt{gr}}(a)$ that is the numerator n of $\mathrm{PrCA}_{\mathcal{F}}^{\texttt{gr}}(a)$. As a final step, algorithm $\mathcal{A}$ returns both n and d.

To prove the statement for *sem* $\in${`ad`, `st`, `co`, `sst`, `pr`, `ids`, `ide`} it suffices to reason analogously to the membership proof for the grounded semantics, by considering that CA^{sem} for those semantics is either in *NP*, or in Θ_2^p, or in $\sum_2^p$.

Specifically, for *sem* $\in${`ad`, `st`, `co`, `sst`, `pr`, `ids`, `ide`}, $\mathrm{PrCA}_{\mathcal{F}}^{sem}(a)$ can be computed by an algorithm $\mathcal{A}'$ which behaves as follows. $\mathcal{A}'$ first computes (in polynomial time) the denominator d of $\mathrm{PrCA}_{\mathcal{F}}^{sem}(a)$. Then $\mathcal{A}'$ invokes a $\#NP$, or a $\#\Theta_2^p$ or a $\#\sum_2^p$ oracle that counts the number of accepting paths of a non-deterministic Turing machine M' defined in the same way of M with the difference that the step (*iii*) of the computation of M is replaced by the invocation of either an NP, a Θ_2^p or a $\sum_2^p$ oracle that checks whether a is an acceptable argument w.r.t. *sem* in the world w obtained as final leaf of the computation tree.
Therefore, since $FP^{\#P} = FP^{\#NP} = FP^{\#\Theta_2^p} = FP^{\#\Sigma_2^p}$, then P-CA^{sem} is in $FP^{\#P}$. □

3.2 Lower bound

To prove that P-CA^{sem} is $FP^{\#P}$-hard we show a reduction from the $\#P$-hard problem $\#P2CNF$, that is the problem of counting the number of satisfying assignments of a CNF formula where each clause consists of exactly 2 positive literals, to P-CA^{sem}.

Definition 4. *Let $\phi = C_1 \wedge C_2 \wedge \ldots C_k$ be a P2CNF, where $X = \{X_1, \ldots, X_n\}$ is the set of its propositional variables.*

We define the PrAAF associated to ϕ ($\mathcal{F}(\phi) = \langle A, P_A, D, P_D\rangle$) as follows:

- $A = \{a, c_1, \ldots, c_k, x_1, \ldots, x_n\}$;
- $D = \{(c_i, a) \mid i \in [1..k]\} \bigcup_{(C_h = X_i \vee X_j) \in \phi} \{(x_i, c_h), (x_j, c_h)\}$;
- $P_A(a) = 1.0$, $\forall i \in [1..k] P_A(c_i) = 1.0$, *and* $\forall j \in [1..n] P_A(x_j) = \frac{1}{2}$;
- $P_D(\delta) = 1.0$ *for each* $\delta \in D$.

Moreover, let γ be a truth assignment for the propositional variables $x_1, \ldots, x_n$. We define as $w(\gamma, \mathcal{F}(\phi))$ the possible world $w = \langle A_w, D_w\rangle \in pw(\mathcal{F}(\phi))$ defined as follows:

- $A_w = A \setminus \{x_i \mid i \in [1..n] \wedge \gamma(x_i) = \texttt{false}\}$;
- $D_w = D \setminus \{(x_i, c_j) \mid i \in [1..n] \wedge j \in [1..k] \wedge \gamma(x_i) = \texttt{false}\}\big)$;

It is easy to see that given a possible world $w \in pw(\mathcal{F}(\phi))$ the probability of w is $I(w) = \frac{1}{2^n}$.

Example 6. Consider the $P2CNF$ formula $\phi = (X_1 \vee X_3) \wedge (X_2 \vee X_3)$. The PrAAf $\mathcal{F}(\phi)$ is reported in Figure 3(a), where probabilities different from 1.0 are reported next to each node and edge. Moreover, consider the truth assignment γ for X_1, X_2 and X_3 such that $\gamma(X_1) = \texttt{true}$, $\gamma(X_2) = \texttt{true}$ and $\gamma(X_3) = \texttt{false}$. The possible world $w(\gamma, \mathcal{F}(\phi))$ is reported in Figure 3(b). □

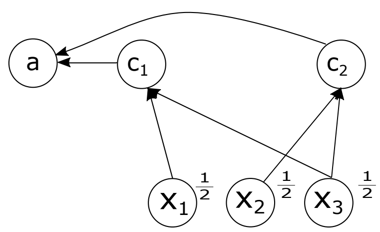

(a)

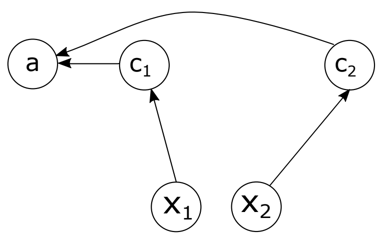

(b)

Fig. 3: Graphical representation of the PrAAF $\mathcal{F}(\phi)$ (a) and the possible world $w(\gamma, \mathcal{F}(\phi))$ (b).

Lemma 1. *Let $\phi = C_1 \wedge C_2 \wedge \ldots C_k$ be a P2CNF, where $X = \{X_1, \ldots, X_n\}$ is the set of its propositional variables.*

For each truth assignment γ for the propositional variables $X_1, \ldots, X_n$ it holds that $\mathrm{CA}^{\boldsymbol{ad}}_{w(\gamma, \mathcal{F}(\phi))}(a)$ *is true iff $\gamma(\phi)$ is true.*

Proof. We first prove that for each truth assignment γ for $X_1, \ldots, X_n$ the fact that $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true implies that $\gamma(\phi)$ is true.

Reasoning by contradiction assume that there exists a truth assignment γ for $X_1, \ldots, X_n$ such that $w = w(\gamma, \mathcal{F}(\phi))$ and $\mathrm{CA}^{\mathtt{ad}}_{w}(a)$ is true while $\gamma(\phi)$ is false.

Since $\mathrm{CA}^{\mathtt{ad}}_{w}(a)$ is true there exists an admissible extension S for w such that $a \in S$. It is easy to see that $S \cap \{c_1, \ldots, c_k\} = \emptyset$, as otherwise S would not be an admissible extension (S would not be conflict-free). Moreover, it is easy to see that there exists a subset $\{x_{i_1}, \ldots, x_{i_h}\}$ of $\{x_1, \ldots, x_n\}$ such that $\{x_{i_1}, \ldots, x_{i_h}\} \subseteq S$ and for each $j \in [1..k]$ there is $l \in [1..h]$ such that:

1. $x_{i_l} \in A_w$, and
2. $(x_{i_l}, c_j) \in D_w$,

as otherwise S would not be an admissible extension as there would be a c_j attacking a which is not attacked by any argument in S.

Since $\gamma(\phi)$ is false it must be the case that there is a $j \in [i..k]$ such that $C_j = X_l \vee X_h$ and $\gamma(X_l) = \gamma(X_h) =$`false`. This implies that $x_l \notin A_w$, $x_h \notin A_w$, $(x_l, c_j) \notin D_w$ and $(x_h, c_j) \notin D_w$, thus contradicting the fact that S is an admissible extension. Hence, it holds that for each truth assignment γ for $x_1, \ldots, x_n$ the fact that $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true implies that $\gamma(\phi)$ is true.

We now prove that for each truth assignment γ for $x_1, \ldots, x_n$ the fact that $\gamma(\phi)$ is true implies that $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true. Reasoning by contradiction assume that there exists a truth assignment γ for $x_1, \ldots, x_n$ such that $w = w(\gamma, \mathcal{F}(\phi))$ and $\mathrm{CA}^{\mathtt{ad}}_{w}(a)$ is false while $\gamma(\phi)$ is true.

Let S be the set of arguments $\{a\} \cup \{x_i \,|\, i \in [1..n] \wedge \gamma(X_i) = \mathtt{true}\}$. We show that S is an admissible extension reasoning by contradiction. Since S is not an admissible extension it must be the case that there is a $j \in [1..k]$ such that $C_j = X_h \vee X_l$ and both $x_h \notin S$ and $x_l \notin S$. Hence, by construction of S it holds that $\gamma(X_h) = \gamma(X_h) =$ `false` which, in turn, implies that $\gamma(\phi) =$ `false`, thus contradicting that γ is a truth assignment γ for $x_1, \ldots, x_n$ such that $\gamma(\phi)$ is true. This implies that for each truth assignment γ for $x_1, \ldots, x_n$ the fact that $\gamma(\phi)$ is true implies that $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true.

Lemma 2. *Let $\phi = C_1 \wedge C_2 \wedge \ldots C_k$ be a P2CNF, where $X = \{X_1, \ldots, X_n\}$ is the set of its propositional variables.*

For a possible world $w \in pw(\mathcal{F}(\phi))$ $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true iff $\mathrm{CA}^{sem}_{w(\gamma,\mathcal{F}(\phi))}(a)$, where sem $\in$\{`st`, `co`, `gr`, `sst`, `pr`, `ids`, `ide`\}, is true.

Proof. The if part is straightforward since stable, complete, grounded, semi-stable, preferred, ideal-set and ideal extensions are admissible extensions. As regards the proof of the only if part, reasoning by contradiction we assume that there is an admissible extension S such that $a \in S$ and there is no extension S' according to *sem*, where *sem* $\in$\{`st`, `co`, `gr`, `sst`, `pr`, `ids`, `ide` \}, such that $a \in S'$.

Consider the set of arguments $\hat{S} = S \cup (Arg(w) \cap \{x_1, \ldots, x_n\})$, and note that $Arg(w) \setminus \hat{S} = \{c_1, \ldots, c_k\}$. It is straightforward to see that, since S is an admissible extension then $\hat{S}$ is an admissible extension. Moreover, it is easy to see that (i) $\hat{S}$ is a stable extension since it defeats each argument in $Arg(w) \setminus \hat{S}$, (ii) $\hat{S}$ is a complete extension since it contains all the argument that are acceptable w.r.t. $\hat{S}$, (iii) $\hat{S}$ is a preferred extension since $\hat{S}$ is an admissible extension and any argument in $Arg(w) \setminus \hat{S}$ attacks an argument in $\hat{S}$, so that for each argument $\alpha \in Arg(w) \setminus \hat{S}$ it holds that $\hat{S} \cup \alpha$ is not conflict-free. Moreover, observe that $\hat{S}$ is the unique preferred extension for w.

Furthermore, since $\hat{S}$ is a complete extension and since removing any argument from $\hat{S}$ make it loose the property that it contains all the arguments that are acceptable w.r.t. it, it holds that $\hat{S}$ is a grounded extension.

Moreover, since $\hat{S} \cup \hat{S}^+ = Arg(w)$ then there is no complete extension $\overline{S}$ such that $\overline{S} \cup \overline{S}^+ \supset \hat{S} \cup \hat{S}^+$. Therefore, since $\hat{S}$ is a complete extension it follows that $\hat{S}$ is a semi-stable extension.

Finally, it is straightforward to see that $\hat{S}$ is an ideal-set and ideal extension for w as it is the unique preferred extension for w.

Hence, for possible world $w \in pw(\mathcal{F}(\phi))$ $\mathrm{CA}^{\mathtt{ad}}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true only if $\mathrm{CA}^{sem}_{w(\gamma,\mathcal{F}(\phi))}(a)$ is true, where $sem \in \{\mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{sst}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}\}$, which completes the proof.

Theorem 2. *For sem* $\in\{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{sst}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}\}$*, it holds that* P-$\mathrm{CA}^{sem}$ *is* $FP^{\#P}$*-hard.*

Proof. Given an instance ϕ of $\#P2CNF$, we construct an instance of P-CA^{sem} by defining a PrAAF $\mathcal{F} = \mathcal{F}(\phi)$ as in Definition 4, and we return $2^n \cdot (PrCA^{sem}_{\mathcal{F}}(a))$.

Since for each possible world $w \in pw(\mathcal{F}(\phi))$ we have that $I(w) = \frac{1}{2^n}$, from Lemmas 1 and 2 it follows that $2^n \cdot (PrCA^{sem}_{\mathcal{F}}(a))$ is the number of satisfying assignments of ϕ.

Hence, as the above described reduction from the $\#P$-hard problem $\#P2CNF$ to P-CA^{sem} is a Cook reduction, this suffices to prove that P-CA^{sem} is $FP^{\#P}$-hard, since a problem is $FP^{\#P}$-hard iff it is $\#P$-hard.

4 RELATED WORK

The main state-of-the-art approaches for handling uncertainty in AAFs by relying on probability theory can be classified in two categories, based on the way they interpret the probabilities of the arguments: those adopting the classical *constellations* approach [27, 12, 32, 7, 9, 25, 30, 19, 20] and those adopting the recent *epistemic* one [33, 29, 28]. As regards the epistemic approach, probabilities and extensions have a different semantics, compared with the constellations approach. Specifically, the probability of an argument represents the degree of belief in the argument (the higher the probability, the more the argument is believed), and a key concept is the "rational" probability distribution, that requires that if the belief in an argument is high, then the belief in the arguments attacked by it is low. In this approach, epistemic extensions are considered rather than Dung's extensions, where an epistemic extension is the set of arguments that are believed to be true to some degree. The interested reader can find a more detailed comparative description of the two categories in [26].

We now focus our attention on the approaches classified as constellations, as the complexity characterization provided in our work refers to this class of PrAAFs. [12] addressed the modeling of jury-based dispute resolutions, and proposed a prAAF where uncertainty is taken into account by specifying probability distribution functions (pdfs) over possible AAFs and showing how an instance of the proposed prAAF can be obtained by specifying a probabilistic assumption-based argumentation framework (introduced by themselves). In the same spirit, [32] defined a prAAF as a pdf over the set of possible AAFs, and introduced a probabilistic version of a fragment of the ASPIC framework [31] that can be used to instantiate the proposed prAAF. [8] addresses the problem of computing all the subgraphs of an AAF in which an argument a belongs to the grounded extension, and [9] extends it by focusing on computing the probability that an argument a belongs to the grounded extension of a probabilistic abstract argumentation framework. In particular, [9] assumes to receive a joint probability distribution over the arguments as input. In fact, providing a joint probability distribution usually means specifying the probability values for all the possible correlations, i.e., $P(a)$, $P(a \wedge b)$, $P(a \wedge b \wedge c)$... and so on. This is analogous to providing the probabilities for all the possible AAFs (since defeats are considered as certain).

Differently from [12] and [32], [30] proposed a prAAF where probabilities are directly associated with arguments and defeats, instead of being associated with possible AAFs, and independence among pairs of arguments/defeats is assumed. After claiming that computing the probability $P^{sem}(S)$ that a set S of arguments is an extension according to *sem* requires exponential time for every semantics, [30] proposed a Monte-Carlo simulation approach to approximate $P^{sem}(S)$. [7, 19, 20] build upon [30]: [7] characterizes different semantics from the approach of [30] in terms of probabilistic logic, as a first step in the direction of creating a uniform logical formalization for all the proposed AAFs of the literature, in order to understand and compare the different approaches. [19, 20], instead, showed that computing $P^{sem}(S)$ is actually tractable for the *admissible* and *stable* semantics, but it is $FP^{\#P}$-complete for other semantics, including *complete*, *grounded*, *preferred* and *ideal-set*. Furthermore, [18, 21] proposed a Monte-Carlo approach to efficiently estimate $P^{sem}(S)$ based on the polynomiality results of [19, 20]. In [30], as well as in [12] and [32], $P^{sem}(S)$ is defined as the sum of the probabilities of the possible AAFs where S is an extension according to semantics *sem*.

In the above-cited works, probability theory is recognized as a fundamental tool to model uncertainty. However, a deeper understanding of the role of probability theory in abstract argumentation was developed only later in [25, 26], where the *justification* and the *premise* perspectives of probabilities of arguments are introduced. According to the former perspective the probability of an argument indicates the probability that it is justified in appearing in the argumentation system. In contrast, the premise perspective views the probability of an argument as the probability that the argument is true based on the degrees to which the premises supporting the argument are believed to be true. Starting from these perspectives, in [26], a formal framework showing the connection among argumentation theory, classical logic, and probability theory was investigated. Furthermore, qualification of attacks is addressed in [27], where an investigation of the meaning of the uncertainty concerning defeats in probabilistic abstract argumentation is provided.

The computational complexity of computing extensions has been thoroughly investigated for classical AAFs [15, 16, 13, 17, 24, 3] with respect to several semantics (a comprehensive overview of argumentation semantics can be found in [1]). In particular, [15] presents a number of results on the complexity of some decision questions for semi-stable semantics, while [13] focuses on ideal semantics; complexity results for preferred semantics can be found in [16].

Complexity results about skeptical and credulous acceptance under admissible, complete, grounded, stable and preferred semantics have been provided in [6, 14, 5], while [13] characterizes the complexity of skeptical and credulous acceptance under ideal and ideal-set.

5 Conclusion and future works

Focusing on the constellations approach with independence proposed in [30], in this paper we characterized the complexity of P-CA^{sem} showing that it is $FP^{\#P}$-complete independently from the adopted semantics. Future work will be devoted to the characterization of the complexity of the problem of computing the probability that an argument is skeptically acceptable w.r.t. a given semantics *sem*. Moreover, another interesting direction for future work is that of finding tractable cases for P-CA^{sem} by identifying structural properties of the argumentation graph that will ensure that P-CA^{sem} is solvable in polynomial time.

References

1. Baroni, P., Caminada, M., Giacomin, M.: An introduction to argumentation semantics. Knowledge Eng. Review 26(4), 365–410 (2011)
2. Baroni, P., Giacomin, M.: Semantics of abstract argument systems. In: Argumentation in Artificial Intelligence, pp. 25–44 (2009)
3. Booth, R., Caminada, M., Dunne, P.E., Podlaszewski, M., Rahwan, I.: Complexity properties of critical sets of arguments. In: Proc. of Computational Models of Argument (COMMA). pp. 173–184 (2014)
4. Caminada, M.: Semi-stable semantics. In: Proc. Int. Conf. Computational Models of Argument (COMMA). pp. 121–130 (2006)
5. Coste-Marquis, S., Devred, C., Marquis, P.: Symmetric argumentation frameworks. In: Proc. of Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU). pp. 317–328 (2005)
6. Dimopoulos, Y., Torres, A.: Graph theoretical structures in logic programs and default theories. Theor. Comput. Sci. 170(1-2), 209–244 (1996)

7. Doder, D., Woltran, S.: Probabilistic argumentation frameworks - A logical approach. In: Proc. Int. Conf. on Scalable Uncertainty Management (SUM). pp. 134–147 (2014)
8. Dondio, P.: Computing the grounded semantics in all the subgraphs of an argumentation framework: An empirical evaluation. In: Proc. Int. Workshop Computational Logic in Multi-Agent Systems (CLIMA). pp. 119–137 (2013)
9. Dondio, P.: Toward a computational analysis of probabilistic argumentation frameworks. Cybernetics and Systems 45(3), 254–278 (2014)
10. Dung, P.M.: 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)
11. Dung, P.M., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. Artif. Intell. 171(10-15), 642–674 (2007)
12. Dung, P.M., Thang, P.M.: Towards (probabilistic) argumentation for jury-based dispute resolution. In: Proc. Int. Conf. Computational Models of Argument (COMMA). pp. 171–182 (2010)
13. Dunne, P.E.: The computational complexity of ideal semantics. Artif. Intell. 173(18) (2009)
14. Dunne, P.E., Bench-Capon, T.J.M.: Coherence in finite argument systems. Artif. Intell. 141(1/2), 187–203 (2002)
15. Dunne, P.E., Caminada, M.: Computational complexity of semi-stable semantics in abstract argumentation frameworks. In: Proc.European Conf. on Logics in Artificial Intelligence (JELIA). pp. 153–165 (2008)
16. Dunne, P.E., Wooldridge, M.: Complexity of abstract argumentation. In: Argumentation in Artificial Intelligence, pp. 85–104 (2009)
17. Dvorák, W., Woltran, S.: Complexity of semi-stable and stage semantics in argumentation frameworks. Inf. Process. Lett. 110(11), 425–430 (2010)
18. Fazzinga, B., Flesca, S., Parisi, F.: Efficiently estimating the probability of extensions in abstract argumentation. In: Proc. Int Conf. on Scalable Uncertainty Management (SUM). pp. 106–119 (2013)
19. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation. In: Proc. Int. Joint Conference on Artificial Intelligence (IJCAI) (2013)
20. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation frameworks. ACM Trans. Comput. Log. (TOCL) 16(3), 22 (2015)
21. Fazzinga, B., Flesca, S., Parisi, F.: On efficiently estimating the probability of extensions in abstract argumentation frameworks. Int. J. Approx. Reasoning 69, 106–132 (2016)
22. Fazzinga, B., Flesca, S., Parisi, F., Pietramala, A.: PARTY: A mobile system for efficiently assessing the probability of extensions in a debate. In: Proc. Int. Conf. on Database and Expert Systems Applications (DEXA). pp. 220–235 (2015)
23. Fazzinga, B., Flesca, S., Parisi, F., Pietramala, A.: Computing or estimating extension's probabilities over structured probabilistic argumentation frameworks. Special Issue on Probabilistic and other Quantitative Approaches to Computational Argumentation of the IfCoLog Journal of Logics and their Applications 3(2), 177–200 (2016)
24. Gaggl, S.A., Woltran, S.: The cf2 argumentation semantics revisited. J. Log. Comput. 23(5), 925–949 (2013)
25. Hunter, A.: Some foundations for probabilistic abstract argumentation. In: Proc. Int. Conf. Computational Models of Argument (COMMA). pp. 117–128 (2012)

26. Hunter, A.: A probabilistic approach to modelling uncertain logical arguments. Int. J. Approx. Reasoning 54(1), 47–81 (2013)
27. Hunter, A.: Probabilistic qualification of attack in abstract argumentation. Int. J. Approx. Reasoning 55(2), 607–638 (2014)
28. Hunter, A., Thimm, M.: Probabilistic argumentation with epistemic extensions. In: Proc. Int. Workshop on Defeasible and Ampliative Reasoning (DARe@ECAI) (2014)
29. Hunter, A., Thimm, M.: Probabilistic argumentation with incomplete information. In: Proc. European Conf. on Artificial Intelligence (ECAI). pp. 1033–1034 (2014)
30. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: Proc. Int. Workshop on Theorie and Applications of Formal Argumentation (TAFA) (2011)
31. Prakken, H.: An abstract framework for argumentation with structured arguments. Argument & Computation 1(2), 93–124 (2010)
32. Rienstra, T.: Towards a probabilistic dung-style argumentation system. In: AT. pp. 138–152 (2012)
33. Thimm, M.: A probabilistic semantics for abstract argumentation. In: Proc. European Conf. on Artificial Intelligence (ECAI). pp. 750–755 (2012)
34. Valiant, L.G.: The complexity of computing the permanent. Theor. Comput. Sci. (TCS) 8, 189–201 (1979)