

Computing extensions' probabilities over probabilistic Bipolar Abstract Argumentation Frameworks

Bettina Fazzinga¹, Sergio Flesca², Filippo Furfaro², and Francesco Scala²

¹ ICAR-CNR, Italy

² DIMES - University of Calabria, Italy

Abstract. *Probabilistic Bipolar Abstract Argumentation Frameworks* (prBAFs), combining the possibility of specifying supports between arguments with a probabilistic modeling of the uncertainty, have been recently considered [34, 35] and the complexity of the problem of computing extensions' probabilities has been characterized [22]. In this paper we deal with the problem of computing extensions' probabilities over prBAFs where the probabilistic events that arguments, supports and defeats occur in the real scenario are assumed to be independent probabilistic events (prBAFs of type IND). Specifically an algorithm for efficiently computing extensions' probabilities under the stable and admissible semantics has been devised and its efficiency has been experimentally validated w.r.t. the exhaustive approach, i.e. the approach consisting in the generation of all the possible scenarios.

1 Introduction

An *abstract argumentation framework* (AAF) represents a dispute as an *argumentation graph* $\langle A, D \rangle$, where A is the set of nodes (called *arguments*) and D is the set of edges (called *defeats* or *attacks*). Herein, an argument is an abstract entity that may attack and/or be attacked by other arguments, where “ a attacks b ” means that argument a rebuts/weakens b . Reasoning on the possible strategies for winning the dispute typically requires looking into the *extensions* of the AAF. An extension S is a set of arguments that satisfies some properties certifying its “strength”, so that a party using the arguments in S has reasonable chances to win the dispute. Different semantics for AAFs (i.e., sets of properties assessing whether a set of arguments is an extension) have proven reasonable, such as *admissible* (*ad*), *stable* (*st*), *preferred* (*pr*), *complete* (*co*), *grounded* (*gr*), *ideal* (*id*) [14, 15, 2], and the complexity of the fundamental problem EXT of verifying whether a set is an extension has been studied under each of these semantics [19, 17].

Since the introduction of AAFs in [14], many variants have been proposed, with the aim of modeling disputes more accurately. Among these, *Bipolar Abstract Argumentation Frameworks* (BAFs) allow *supports*, besides attacks, to be specified between arguments. Specifically, two alternative formal semantics of support have been introduced: in [6], the support is a generic “inverse” of the notion of attack (“*abstract semantics*”: “ a supports b ” means that a strengthens the validity of a), while, in [4], it is viewed as a “deductive” correlation between arguments (“*deductive semantics*”: if a supports b , the acceptance of a implies the acceptance of b). The various extensions' semantics defined

for AAFs have been shown to have a natural counterpart over BAFs, after noticing that combining attacks with supports (of any semantics) generates “implicit” attacks (see Example 1).

Example 1. The graph in Figure 1 is a BAF with six arguments a, b, c, d, e, f . The dashed and the standard arrows denote supports and attacks, respectively. The co-existence of supports and attacks entails the existence of implicit attacks. For instance, under both the abstract and deductive semantics, the fact that a strengthens b and b attacks c implicitly says that a attacks c . This kind of implicit attack is often called “*supported attack*”. If the deductive semantics is adopted, there are other forms of implicit attacks. For instance, since a supports b and e attacks b , there is an implicit attack from e to a . Otherwise, a would be acceptable while b would be not, thus contradicting the deductive support from a to b .

Other variants of AAFs that, owing to their practical impact, have gained interest from the research community are those addressing the representation of uncertainty. In this regard, *probabilistic AAFs* (prAAFs) are a popular paradigm, and in particular those following the *constellation approach*. Here, the dispute is modeled as a set of possible scenarios, each consisting of a standard AAF (called *possible AAF*) associated with a probability of representing all and only the arguments and attacks actually occurring in the dispute. In particular, two main paradigms have been adopted for specifying the probability distribution function (pdf), called EX and IND. In the general case, the *extensive* form EX is used, where the composition of each possible AAF must be explicitly specified along with its probability. Otherwise, when *independence* between arguments/attacks is assumed, the form IND can be used, where the possible scenarios and their probabilities are represented compactly and implicitly by specifying the marginal probabilities of the arguments and attacks.

Recently, for both EX and IND, the complexity of the probabilistic counterpart P-EXT of EXT (asking for the probability that a set of arguments is an extension) has been characterized for IND in [23] for prAFFs and in [22] for prBAFs. Interestingly, when considering IND and the stable or the admissible semantics, moving from prAAFs to prBAFs makes P-EXT intractable ($FP^{\#P}$ -complete).

In this paper, we consider the *probabilistic Bipolar Argumentation Framework* (prBAF), where the bipolarity of BAFs is combined with the probabilistic modeling of the uncertainty of prAAFs under the paradigm IND, and provide an efficient algorithm for solving P-EXT in the cases that the stable or the admissible semantics are considered. The algorithm is shown to be sound and its efficiency has been experimentally validated.

Related Work. [6] first introduced BAFs, where supports have the general “abstract” semantics of positive interactions between arguments. Later, three more specific interpretations for supports have been proposed: [4], [32] and [33] introduced the *deductive*,




Fig. 1: A bipolar abstract argumentation framework

necessary, and *evidential* semantics for the support relation, respectively. In this paper, we focus on the abstract and deductive semantics, but our results also hold for necessary supports (as shown in [8], they are dual to deductive ones). [8] reviews the four different semantics for supports, and discusses the similarities and differences among these interpretations. [9] introduces a more general framework that incorporates attacks, supports and a preference relation to decide between conflicting arguments. In [30], subarguments in AAF have been introduced, that in [10] have been shown to be closely related with the necessary support. Other related works are [5, 38] where, although supports are not mentioned, similar dependencies have been considered. A detailed survey over BAFs can be found in [10].

Regarding uncertainty in AAFs, the approaches based on probability theory can be classified in two categories: those adopting the classical *constellations* approach [26, 16, 36, 12, 13, 24, 29, 20, 21] and those adopting the recent *epistemic* one [37, 28, 27]. The former category has the two sub-categories EX [16, 36, 13] and IND [12, 29, 20, 21], described in the paper. The interested reader can find a more detailed comparative description of the two categories in [25]. Furthermore, many proposals have been made where uncertainty is represented by exploiting weights or preferences on arguments and/or defeats [3, 1, 31, 18, 11]. Although the approaches based on weights, preferences, or probabilities to model uncertainty have proved effective in different contexts, there is no common agreement on what kind of approach should be used in general. In this regard, [24, 25] observed that the probability-based approaches may take advantage from relying on a well-established and well-founded theory, whereas the approaches based on weights or preferences do not.

As regards *probabilistic Bipolar Argumentation Framework* (prBAF) they have been recently introduced [34, 35] and the computational complexity of the problem of computing extensions' probabilities has been characterized [22].

2 Preliminaries

We assume that the reader is familiar with the notions of *Abstract Argumentation Framework* (AAF), extension, and acceptability of arguments. We now review *Bipolar Abstract Argumentation Frameworks* (BAFs) and the concepts of support, attack and defense, along with the most popular extensions' semantics over BAFs [6].

2.1 Bipolar abstract argumentation frameworks

Definition 1. [BAF] A bipolar abstract argumentation framework (BAF) is a tuple $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, where \mathcal{A} is a set of arguments, $\mathcal{R}_a \subseteq \mathcal{A} \times \mathcal{A}$ is a defeat/attack relation and $\mathcal{R}_s \subseteq \mathcal{A} \times \mathcal{A}$ is a support relation.

In the first proposal of BAF [6], supports are given an *abstract semantics*, that is the opposite of the traditional semantics of attack, inherited from classical AAFs. This was shown to make the combination of supports and attacks imply the so-called *supported attacks*³.

³ [6] discussed also *indirect* implicit attacks. W.l.o.g., as in [7], we disregard them, as their presence would not affect our results.

Definition 2. [Supported attack] Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ be a BAF, and $a, b \in \mathcal{A}$. There is a supported attack from a to b iff there is a sequence $a_1 \mathcal{R}_1 \dots \mathcal{R}_{n-1} a_n$, with $n \geq 2$, where $a_1 = a$, $a_n = b$, $\forall i \in [1..n-2] \mathcal{R}_i = \mathcal{R}_s$, and $\mathcal{R}_{n-1} = \mathcal{R}_a$.

Example 2. In the BAF in Figure 1, there is a supported attack from a to c . Also the three direct attacks (e, b) , (b, c) and (f, d) are special cases of supported attack.

Besides the abstract, other semantics have been proposed for supports (see *Related Work*). In particular, we consider the well-established *deductive semantics*, first proposed in [4]. Here, “ a supports b ” is interpreted as a strong correlation between a and b , meaning that if a is acceptable, then b is acceptable too. As observed in [8], under this semantics, a new form of implicit attack, called *d-attack*, must be considered.

Definition 3. [d-attack] Given a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ and $a, b \in \mathcal{A}$, there is a d-attack from a to b iff

- $a \mathcal{R}_a b$, or
- there is an argument a' such that there is a path from a to a' consisting of only support edges, and a' attacks b , or
- there is an argument a' such that there is a path from b to a' consisting of only support edges, and a attacks a' .

Example below shows that d-attacks include supported attacks, but can be also of the form of *supermediated attacks* (described by the last point in Definition 3).

Example 3. Under the deductive semantics for supports, in the BAF of Figure 1, it is easy to see that the supported attacks reported in Example 2 are d-attacks. Further d-attacks are the supermediated attacks from e to a , and from f to c , that are not supported attacks.

In order to analyze what changes when moving from one semantics of supports to the other (or, equivalently, from one form of implicit attacks to the other), we partition BAFs into two classes: *s-BAFs* and *d-BAFs*, where only supported attacks and d-attacks are considered, respectively. From now on, we assume the presence of a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, and, when needed, we will specify whether \mathcal{F} is an s- or a d-BAF.

The concepts of *support*, *attack* and *defense from sets of arguments* are mandatory to define the extensions over BAFs.

Definition 4. [Set-support] A set $S \subseteq \mathcal{A}$ set-supports an argument $a \in \mathcal{A}$ iff there is an argument $a' \in S$ such that there is a path from a' to a consisting of only support edges.

Definition 5. [Set-attack] Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ be an s-BAF (resp., d-BAF). A set $S \subseteq \mathcal{A}$ set-attacks $a \in \mathcal{A}$ iff there is a supported attack (resp., d-attack) from some $b \in S$ to a .

Definition 6. [Set-defense] A set $S \subseteq \mathcal{A}$ set-defends an argument $a \in \mathcal{A}$ iff, $\forall b \in \mathcal{A}$, if $\{b\}$ set-attacks a then $\exists c \in S$ such that $\{c\}$ set-attacks b .

Example 4. Consider the BAF \mathcal{F} in Figure 1. Independently from supports' semantics, $\{a, e\}$ both set-supports and set-attacks b , and also set-attacks c . If \mathcal{F} is an s-BAF, then $\{a, e\}$ does not set-defend c , since there is a supported attack from a to c , and no attack from e to a . Observe that $\{a, e\}$ does not set-defend c if \mathcal{F} is a d-BAF either, since, although e d-attacks a , no argument in $\{a, e\}$ d-attacks f , that in turn d-attacks c .

2.2 Semantics

We first recall the notions of conflict-freeness and safety.

Definition 7. [Conflict-free and safe sets of arguments] *A set of arguments $S \subseteq \mathcal{A}$ is:*

- conflict-free iff $\nexists a, b \in S$ such that $\{a\}$ set-attacks b ;
- safe iff $\nexists b \in \mathcal{A}$ such that S set-attacks b and either S set-supports b or $b \in S$.

Example 5. If the BAF \mathcal{F} in Figure 1 is an s-BAF, both $\{a, e\}$, and $\{f, c\}$ are conflict-free but not safe, while both $\{a, b, f\}$ and $\{a, b, d\}$ are conflict-free and safe. If \mathcal{F} is a d-BAF, both $\{a, e\}$, and $\{f, c\}$ are not conflict-free, while both $\{a, b, f\}$ and $\{a, b, d\}$ are still conflict-free and safe.

All the most popular semantics of extensions of “standard” AAFs have been extended to the case of BAFs [6]. We start with the stable semantics.

Definition 8. [Stable extension] *A set of arguments $S \subseteq \mathcal{A}$ is a stable extension iff S is conflict-free and $\forall a \in \mathcal{A} \setminus S$ it holds that S set-attacks a .*

The presence of supports and the fact that, in BAFs, conflict-freeness and safety do not coincide (while they do in AAFs) is at the basis of the fact that, for some AAF's semantics, different variants are considered when moving to BAFs. This is the case of the admissible semantics.

Definition 9. [Admissible extension] *A set $S \subseteq \mathcal{A}$ is*

- a d-admissible extension iff S is conflict-free and set-defends all of its arguments;
- an s-admissible extension iff S is safe and set-defends all of its arguments;
- a c-admissible extension iff S is conflict-free, closed for \mathcal{R}_s and set-defends all of its arguments.

In turn, the other semantics subsuming the admissible one are defined as follows. A set $S \subseteq \mathcal{A}$ is said to be:

- a d-complete (resp. s-complete, c-complete) extension iff S is d-admissible (resp., s-admissible, c-admissible) and S contains all the arguments set-defended by S ;
- a d-grounded (resp. s-grounded, c-grounded) extension iff S is a minimal (w.r.t. \subseteq) d-complete (resp. s-complete, c-complete) extension;
- a d-preferred (resp. s-preferred, c-preferred) extension iff S is a maximal (w.r.t. \subseteq) d-complete (resp. s-complete, c-complete) extension;

- a *d-ideal* (resp. *s-ideal*, *c-ideal*) extension iff S is a maximal (w.r.t. \subseteq) *d-admissible* (resp. *s-admissible*, *c-admissible*) extension and S is contained in every *d-preferred* (resp. *s-preferred*, *c-preferred*) extension.

We denote the set $\{d-ad, s-ad, c-ad, st, d-co, s-co, c-co, d-gr, c-gr, c-gr, d-pr, s-pr, c-pr, d-id, s-id, c-id\}$ consisting of the above semantics as SEM (herein, *st* means stable, *d-ad* *d-admissible*, *s-ad* *s-admissible*, and so on).

Example 6. Consider the BAF in Figure 1. Both $\{a, b, f\}$ and $\{a, b, d\}$, although conflict-free and safe, are not *d-ad* extensions in both the cases of s-BAF and d-BAF (since b is not set-defended). Furthermore, for the s-BAF case, we have: both $\{a, f\}$ and $\{e, f\}$ are *s-ad*, *s-gr* and *s-pr* extensions, $\{a, e, f\}$ is a *st*, *d-ad*, *d-gr*, *d-pr*, and *d-id* extension, $\{f\}$ is an *s-id* extension, $\{e, f\}$ is a *c-pr*, *c-gr* and *c-id* extension.

For the d-BAF case we have: $\{e, f\}$ is the unique stable extension, that is also *c-preferred*, *c-grounded* and *c-ideal*.

The fundamental problem of verifying whether a set S of arguments is an extension over a given BAF under a semantics $sem \in SEM$ will be denoted as $EXT^{sem}(S)$. Basically, solving an instance of $EXT^{sem}(S)$ means checking whether a set of arguments is a reasonable strategy in the dispute, where the meaning of “reasonable” is encoded in the semantics.

Given a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, a set $S \subseteq A$, and a semantics $sem \in SEM$, we define the boolean function $ext(\alpha, sem, S)$ returning *true* iff S is an extension under sem .

3 Probabilistic BAFs (prBAFs)

We now consider the extension of BAFs where uncertainty is addressed and modeled probabilistically as in “traditional” *probabilistic Abstract Argumentation Frameworks* - prAAFs (in particular, we refer to prAAFs employing the constellation approach recalled in the introduction and adopting the IND approach for specifying pdfs).

A *probabilistic BAF* (prBAF) \mathcal{F} of type IND is a tuple $\langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$ where $A = \{a_1, \dots, a_m\}$, $\mathcal{R}_a = \{\delta_1, \dots, \delta_n\}$ and $\mathcal{R}_s = \{\sigma_1, \dots, \sigma_k\}$ are the sets of arguments, attacks and supports, respectively, and $\mathcal{P}_A = \{P(a_1), \dots, P(a_m)\}$, $\mathcal{P}_R = \{P(\delta_1), \dots, P(\delta_n), P(\sigma_1), \dots, P(\sigma_k)\}$ are their marginal probabilities.

A prBAF \mathcal{F} is used to represent a set of *possible BAFs*, that is the alternative cases of dispute that may occur, and their probabilities. More in detail $PS = \{\alpha = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle \mid \mathcal{A}' \subseteq \mathcal{A} \wedge \mathcal{R}'_a \subseteq (\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}_a \wedge \mathcal{R}'_s \subseteq (\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}_s\}$ is the set of possible BAFs represented by \mathcal{F} and the pdf P over the possible scenarios that is implied by the independence assumption and the marginal probabilities $\mathcal{P}_A, \mathcal{P}_R$ is as follows. For each possible BAF $\alpha' = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle$, the probability $P(\alpha')$ is:

$$\begin{aligned}
P(\alpha') = & \prod_{a \in \mathcal{A}'} P(a) \times \prod_{a \in \mathcal{A} \setminus \mathcal{A}'} (1 - P(a)) \times \\
& \prod_{\delta \in \mathcal{R}'_a} P(\delta) \times \prod_{\delta \in (\mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}')) \setminus \mathcal{R}'_a} (1 - P(\delta)) \times \\
& \prod_{\sigma \in \mathcal{R}'_s} P(\sigma) \times \prod_{\sigma \in (\mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}')) \setminus \mathcal{R}'_s} (1 - P(\sigma)).
\end{aligned} \tag{1}$$

The size of a prBAF of type IND is $O(|\mathcal{A}| + |\mathcal{R}_a| + |\mathcal{R}_s| + |\mathcal{P}_A| + |\mathcal{P}_R|)$.

Example 7. Consider a prBAF \mathcal{F}' of form IND, where \mathcal{A} , \mathcal{R}_a and \mathcal{R}_s are those of Figure 1, and $\mathcal{P}_\mathcal{A}$ and $\mathcal{P}_\mathcal{R}$ are the following: $P(a) = P(b) = P(c) = P(d) = P(f) = 1$, $P(e) = 0.5$, $P(e, b) = 0.5$ and the probabilities of the other supports and attacks are equal to 1. We have three possible scenarios: $\alpha_1 = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, $\alpha_2 = \langle \mathcal{A}, \mathcal{R}_a \setminus \{(e, b)\}, \mathcal{R}_s \rangle$, $\alpha_3 = \langle \mathcal{A} \setminus \{e\}, \mathcal{R}_a \setminus \{(e, b)\}, \mathcal{R}_s \rangle$, whose probabilities are: $P(\alpha_1) = 0.25$, $P(\alpha_2) = 0.25$, $P(\alpha_3) = 0.5$.

In what follows, given a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_\mathcal{A}, \mathcal{P}_\mathcal{R} \rangle$, we denote as $\mathcal{F}.\alpha = \alpha_1, \dots, \alpha_m$ the possible BAFs that are assigned non-zero probability by \mathcal{P} , and as $\mathcal{F}.\mathbf{P} = P(\alpha_1), \dots, P(\alpha_m)$ their probabilities.

The probabilistic versions of the two sub-classes s-BAF and d-BAF will be called *s-prBAF* and *d-prBAF*, respectively.

When switching to the probabilistic setting, the decision problem $\text{EXT}^{sem}(S)$ makes no sense, since a number of different scenarios are possible, and a set of arguments can be an extension in some scenarios, but not in others. Thus, the most natural “translation” of the problem of examining the “reasonability” of a set of arguments S becomes the functional problem $\text{P-EXT}^{sem}(S)$ of evaluating the probability that S is an extension, according to the following definition.

Definition 10 (P-EXT $^{sem}(S)$ and $P^{sem}(S)$). *Given a prBAF \mathcal{F} , a set S of arguments, and a semantics $sem \in SEM$, P-EXT $^{sem}(S)$ is the problem of computing the probability $P_{\mathcal{F}}^{sem}(S)$ that S is an extension under sem , i.e.*

$$P_{\mathcal{F}}^{sem}(S) = \sum_{\alpha \in \mathcal{F}.\alpha \wedge ext(\alpha, sem, S)} \mathcal{F}.\mathbf{P}(\alpha) \quad (2)$$

Example 8. Continuing examples 6 and 7, we now compute the probability that $S = \{a, e\}$ is d-admissible in both the s- and d- prBAF cases.

Case s-prBAF: S is d-admissible in both α_1 and α_2 (as e is missing in α_3), thus $P^{d-ad}(S) = P(\alpha_1) + P(\alpha_2) = 0.5$.

Case d-prBAF: S is d-admissible only in α_2 , as in α_1 e d-attacks a and in α_3 e is missing, thus we have $P^{d-ad}(S) = P(\alpha_2) = 0.25$.

4 An algorithm for computing $P_{\mathcal{F}}^{sem}(S)$

We now provide our main contribution, that is the definition of an algorithm for efficiently computing the probability that a set of arguments is an extension according to a semantics $sem \in \{st, d-ad, s-ad, c-ad\}$.

The algorithm is based on the following results, provided in [21], that hold for traditional prAAF of type IND. A PrAAF of type IND correspond to a prBAF of type IND where the support relation is empty.

Fact 1 [21] *Given a prAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \emptyset, \mathcal{P}_\mathcal{A}, \mathcal{P}_\mathcal{R} \rangle$ and a set $S \subseteq \mathcal{A}$ of arguments, the probability that S is an admissible extension in \mathcal{F} is equal to $P_1(S) \cdot P_2(S) \cdot P_3(S)$, where⁴:*

⁴ Note that an empty product evaluates to 1.

- $P_1(S) = \prod_{a \in S} \mathcal{P}_{\mathcal{A}}(a)$,
- $P_2(S) = \prod_{\substack{(a,b) \in \mathcal{R}_a \\ \wedge a \in S \\ \wedge b \in S}} (1 - \mathcal{P}_{\mathcal{R}}(\langle a, b \rangle))$, and
- $P_3(S) = \prod_{d \in \mathcal{A} \setminus S} (P_{31}(S, d) + P_{32}(S, d) + P_{33}(S, d))$, where:
 - $P_{31}(S, d) = 1 - \mathcal{P}_{\mathcal{A}}(d)$,
 - $P_{32}(S, d) = \mathcal{P}_{\mathcal{A}}(d) \times \prod_{\substack{(d,b) \in \mathcal{R}_a \\ \wedge b \in S}} (1 - \mathcal{P}_{\mathcal{R}}(\langle d, b \rangle))$,
 - $P_{33}(S, d) = \mathcal{P}_{\mathcal{A}}(d) \times \left(1 - \prod_{\substack{(d,b) \in \mathcal{R}_a \\ \wedge b \in S}} (1 - \mathcal{P}_{\mathcal{R}}(\langle d, b \rangle))\right) \times$
 $\left(1 - \prod_{\substack{(a,d) \in \mathcal{R}_a \\ \wedge a \in S}} (1 - \mathcal{P}_{\mathcal{R}}(\langle a, d \rangle))\right)$

Fact 2 [21] Given a prAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \emptyset, \mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{R}} \rangle$ and a set $S \subseteq \mathcal{A}$ of arguments, the probability that S is a stable extension in \mathcal{F} is equal to $P_1(S) \cdot P_2(S) \cdot P_3(S)$, where $P_1(S)$ and $P_2(S)$ are defined as in Fact 1 and

$$P_3(S) = \prod_{d \in \mathcal{A} \setminus S} (P_{31}(S, d) + P_{32}(S, d))$$

where:

- $P_{31}(S, d) = 1 - \mathcal{P}_{\mathcal{A}}(d)$, and
- $P_{32}(S, d) = \mathcal{P}_{\mathcal{A}}(d) \times \left(1 - \prod_{\substack{(a,d) \in \mathcal{R}_a \\ \wedge a \in S}} (1 - \mathcal{P}_{\mathcal{R}}(\langle a, d \rangle))\right)$

Before defining the algorithm we introduce some preliminary notations used in its definition. Formally, given a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{R}} \rangle$, we consider the sets:

- $\mathcal{F}.\mathcal{A}_e = \{a \mid \exists \langle a, b \rangle \in \mathcal{R}_s \vee \exists \langle b, a \rangle \in \mathcal{R}_s\}$ (called set of *supp-arguments*, as they are those involved in supports), and
- $\mathcal{F}.\mathcal{R}_e = \{\langle a, b \rangle \in (\mathcal{R}_a \cup \mathcal{R}_s) \mid (a \in \mathcal{A}_e \vee b \in \mathcal{A}_e)\}$ (called set of *supp-attacks and supports*, as they are attacks/supports incident to supp-arguments).

The algorithm evaluation strategy is based on the notions of *contraction* and *completion*. A *contraction* for a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{R}} \rangle$ is a prBAF $\mathcal{F}^* = \langle \mathcal{A}^*, \mathcal{R}_a^*, \mathcal{R}_s^*, \mathcal{P}_{\mathcal{A}}^*, \mathcal{P}_{\mathcal{R}}^* \rangle$ where:

- $\mathcal{A}^* \subseteq \mathcal{A}$ and $\mathcal{A} \setminus \mathcal{A}^* \subseteq \mathcal{F}.\mathcal{A}_e$;
- $\mathcal{R}_a^* \subseteq \mathcal{R}_a \cap (\mathcal{A}^* \times \mathcal{A}^*)$ and $\mathcal{R}_a \setminus \mathcal{R}_a^* \subseteq \mathcal{F}.\mathcal{R}_e$;
- $\mathcal{R}_s^* \subseteq \mathcal{R}_s \cap (\mathcal{A}^* \times \mathcal{A}^*)$;
- $\mathcal{P}_{\mathcal{A}}^*(a) = 1$ if $a \in \mathcal{F}.\mathcal{A}_e$ and $\mathcal{P}_{\mathcal{A}}(a) = \mathcal{P}_{\mathcal{A}}(a)$ otherwise;
- $\mathcal{P}_{\mathcal{R}}^*(\langle a, b \rangle) = 1$ if $a \in \mathcal{F}.\mathcal{A}_e \vee b \in \mathcal{F}.\mathcal{A}_e$, and $\mathcal{P}_{\mathcal{R}}^*(\langle a, b \rangle) = \mathcal{P}_{\mathcal{R}}(\langle a, b \rangle)$ otherwise.

Basically, \mathcal{F}^* 's supports are a subset of \mathcal{F} 's, and arguments (resp., attacks) are a subset of \mathcal{F} 's containing at least the non supp-arguments (resp., the non supp-attacks). Then, the probabilities are copied from those specified in \mathcal{F} , except for those over supp-arguments and attacks, that are overwritten with 1.

$\mathcal{E}(\mathcal{F})$ will denote the set of possible contractions of \mathcal{F} . For $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$, we define the probability of \mathcal{F}^* given \mathcal{F} as:

$$P(\mathcal{F}^*|\mathcal{F}) = \prod_{a \in \mathcal{A}_e} P_{\mathcal{F}^*}(a) \times \prod_{\delta \in \mathcal{R}_e \cap (\mathcal{A}^* \times \mathcal{A}^*)} P_{\mathcal{F}^*}(\delta), \quad (3)$$

where:

- if $a \in \mathcal{A}^*$, $P_{\mathcal{F}^*}(a) = \mathcal{P}_{\mathcal{A}}(a)$; else, $P_{\mathcal{F}^*}(a) = 1 - \mathcal{P}_{\mathcal{A}}(a)$;
- if $\delta \in \mathcal{R}_a^* \cup \mathcal{R}_s^*$, $P_{\mathcal{F}^*}(\delta) = \mathcal{P}_{\mathcal{R}}(\delta)$; else, $P_{\mathcal{F}^*}(\delta) = 1 - \mathcal{P}_{\mathcal{R}}(\delta)$.

As regards *completions*, their definition uses the function $cert(\mathcal{F})$, returning the BAF consisting of all and only the arguments in $\mathcal{F}.\mathcal{A}_e$ plus the arguments in \mathcal{F} involved in the attacks/supports in $\mathcal{F}.\mathcal{R}_e$ and the attacks/supports of \mathcal{F} that appear in $\mathcal{F}.\mathcal{R}_e$. More formally, given a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{R}} \rangle$, $cert(\mathcal{F})$ is the BAF $\langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle$ such that:

- $\mathcal{A} = \mathcal{F}.\mathcal{A}_e \cup \{a \mid a \in \mathcal{A} \wedge \exists b \in \mathcal{A} s.t. \langle a, b \rangle \in \mathcal{F}.\mathcal{R}_e \vee \langle b, a \rangle \in \mathcal{F}.\mathcal{R}_e\}$,
- $\mathcal{R}'_a = \mathcal{R}_a \cap \mathcal{F}.\mathcal{R}_e$,
- $\mathcal{R}'_s = \mathcal{R}_s \cap \mathcal{F}.\mathcal{R}_e$.

Thus, the *completion* of \mathcal{F} is the prBAF $compl(\mathcal{F}) = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s, \mathcal{P}'_{\mathcal{A}}, \mathcal{P}'_{\mathcal{R}} \rangle$ where:

- $\mathcal{A}' = \mathcal{A}$ and $\mathcal{R}'_s = \mathcal{R}_s$;
- $\mathcal{R}'_a = \mathcal{R}_a \cup R'$, where R' consists of the s- or the d- attacks of $cert(\mathcal{F})$, depending on whether \mathcal{F} is an s- or a d- prBAF;
- $\forall a \in \mathcal{A}, \mathcal{P}'_{\mathcal{A}}(a) = \mathcal{P}_{\mathcal{A}}(a)$;
- $\forall \delta \in \mathcal{R}'_a$, if $\delta \in R'$ then $\mathcal{P}'_{\mathcal{R}}(\delta) = 1$, else $\mathcal{P}'_{\mathcal{R}}(\delta) = \mathcal{P}_{\mathcal{R}}(\delta)$.
- $\forall \delta \in \mathcal{R}'_s, \mathcal{P}'_{\mathcal{R}}(\delta) = \mathcal{P}_{\mathcal{R}}(\delta)$.

We now define Algorithm 1 that computes the probability that a set of arguments S is an extension for \mathcal{F} according to the semantics $sem \in \{st, d-ad, s-ad, c-ad\}$ for both s- and d-prBAFs by iterating over $\mathcal{E}(\mathcal{F})$.

Algorithm 1 first computes the sets $\mathcal{F}.\mathcal{A}_e$ and $\mathcal{F}.\mathcal{R}_e$ of supp-arguments and supp-attacks and supports of \mathcal{F} (Lines 2-3) and it initializes Pr to 0. Then it iterates over the possible contractions of \mathcal{F} by iterating over the subsets of $\mathcal{F}.\mathcal{A}_e$ (Line 4) and then for each subset \mathcal{A}'_e of $\mathcal{F}.\mathcal{A}_e$ iterating over the subsets of $\mathcal{R}_e \cap (\mathcal{A}'_e \times \mathcal{A}'_e)$ (Line 5).

The contraction \mathcal{F}^* of \mathcal{F} corresponding to the sets \mathcal{A}'_e and \mathcal{R}'_e is generated by calling function $contract$ (Line 6) and its probability $P(\mathcal{F}^*|\mathcal{F})$ is computed using Equation 3 (Line 7). Then the completion \mathcal{F}' of \mathcal{F}^* is computed as by calling function $complete$ (Line 8).

Then the variable Pr' is computed according to the following definition (Lines 9-19):

- $Pr' = 0.0$ if $sem = s-ad$ and S is not safe in $cert(\mathcal{F}')$,
- $Pr' = 0.0$ if $sem = s-ad$ and S is not closed for supports over $cert(\mathcal{F}')$,
- $Pr' = 1.0$, otherwise.

Next, the probability \overline{Pr} that S is an admissible/stable extension in the prAAF obtained removing supports by \mathcal{F}' ($\overline{\mathcal{F}}$) is computed according to the formulas reported in Facts 1 and 2 (Lines 20- 26). Specifically function $computePrAAF$ is responsible for computing the probability \overline{Pr} that S is an admissible/stable extension in $\overline{\mathcal{F}}$ and \overline{Pr} is added to the probability Pr . Finally Pr is returned.

Algorithm 1 Computing $P_{\mathcal{F}}^{sem}(S)$ by enumerating contractions

Require: A prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$
 A set $S \subseteq A$
 A semantics $sem \in \{st, d-ad, s-ad, c-ad\}$

Ensure: $P_{\mathcal{F}}^{sem}(S)$

- 1: $Pr = 0.0$
- 2: $\mathcal{A}_e = \text{compute}\mathcal{A}_e(\mathcal{F})$
- 3: $\mathcal{R}_e = \text{compute}\mathcal{R}_e(\mathcal{F}, \mathcal{A}_e)$
- 4: **for** $\mathcal{A}_e^* \subseteq \mathcal{A}_e$ **do**
- 5: **for** $\mathcal{R}_e^* \subseteq \mathcal{R}_e \cap (\mathcal{A}_e^* \times \mathcal{A}_e^*)$ **do**
- 6: $\mathcal{F}^* = \text{contract}(\mathcal{F}, \mathcal{A}_e^*, \mathcal{R}_e^*, \mathcal{A}_e, \mathcal{R}_e)$
- 7: $Pr^* = \prod_{a \in \mathcal{A}_e} P_{\mathcal{F}^*}(a) \times \prod_{\delta \in \mathcal{R}_e \cap (\mathcal{A}_e^* \times \mathcal{A}_e^*)} P_{\mathcal{F}^*}(\delta)$
- 8: $\mathcal{F}' = \text{complete}(\mathcal{F}^*, \mathcal{A}_e^*, \mathcal{R}_e^*)$
- 9: **if** $sem = s-ad$ **then**
- 10: **if** not safe($\mathcal{F}', S, \mathcal{A}_e^*, \mathcal{R}_e^*$) **then**
- 11: $Pr' = 0.0$
- 12: **end if**
- 13: **else if** $sem = c-ad$ **then**
- 14: **if** not supclosed(\mathcal{F}', S) **then**
- 15: $Pr' = 0.0$
- 16: **end if**
- 17: **else**
- 18: $Pr' = 1.0$
- 19: **end if**
- 20: $\overline{\mathcal{F}} = \text{toPrAAF}(\mathcal{F}^*)$
- 21: **if** $sem = st$ **then**
- 22: $aa\,f\,sem = st$
- 23: **else**
- 24: $aa\,f\,sem = ad$
- 25: **end if**
- 26: $\overline{Pr} = \text{computePrAAF}(\overline{\mathcal{F}}, S, aa\,f\,sem)$
- 27: $Pr = Pr + Pr^* \times Pr' \times \overline{Pr}$
- 28: **end for**
- 29: **end for**
- 30: **return** Pr

4.1 Correctness of Algorithm 1

In this section we show that Algorithm 1 is sound and characterize its computational complexity. First we introduce a lemma (which straightforwardly follows from the definition of $P(\mathcal{F}^*|\mathcal{F})$) that allows for decomposing the evaluation of $P_{\mathcal{F}}^{sem}(S)$ into evaluating $P_{\mathcal{F}^*}^{sem}(S)$ over each contraction \mathcal{F}^* .

Lemma 1. *Let \mathcal{F} be a prBAF and S a set of its arguments. For $sem \in \{d-ad, s-ad, c-ad, st\}$, it holds that $P_{\mathcal{F}}^{sem}(S) = \sum_{\mathcal{F}^* \in \mathcal{E}(\mathcal{F})} P(\mathcal{F}^*|\mathcal{F}) \times P_{\mathcal{F}^*}^{sem}(S)$.*

The following lemma state that the method for computing $P_{\mathcal{F}^*}^{sem}(S)$ for each $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$ used in Algorithm 1 is correct. Indeed, it states that $P_{\mathcal{F}^*}^{sem}(S)$ can be computed by taking the prAAF $\overline{\mathcal{F}}$ obtained by removing the supports from the completion of

\mathcal{F}^* , and then using over $\overline{\mathcal{F}}$ any state-of-the-art algorithm for computing the extensions' probabilities over "traditional" prBAFs (e.g. using the formulas reported in Facts 1 and 2).

Lemma 2. *Let \mathcal{F} be a prBAF, \mathcal{F}^* a contraction for \mathcal{F} , $\overline{\mathcal{F}}$ the prAAF obtained from $\text{compl}(\mathcal{F}^*)$ by removing the supports, and S a set of arguments of \mathcal{F}^* . Then:*

- For $sem \in \{d\text{-ad}, st\}$ $P_{\mathcal{F}^*}^{sem}(S) = P_{\overline{\mathcal{F}}}^{sem'}(S)$, where $sem' \in \{ad, st\}$ respectively;
- $P_{\mathcal{F}^*}^{s\text{-ad}}(S) = 0$ if S is not safe over $\text{cert}(\text{compl}(\mathcal{F}^*))$; otherwise, $P_{\mathcal{F}^*}^{s\text{-ad}}(S) = P_{\overline{\mathcal{F}}}^{ad}(S)$;
- $P_{\mathcal{F}^*}^{c\text{-ad}}(S) = 0$ if S is not closed for \mathcal{R}_s over $\text{cert}(\text{compl}(\mathcal{F}^*))$; otherwise, $P_{\mathcal{F}^*}^{c\text{-ad}}(S) = P_{\overline{\mathcal{F}}}^{ad}(S)$.

Proof (Sketch). The statement can be proved by exploiting the fact that, since supports and attacks in contractions are certain, safeness and closure for \mathcal{R}_s hold over \mathcal{F}^* iff they hold over $\text{cert}(\text{compl}(\mathcal{F}^*))$. The detailed proof is omitted for space reasons. \square

Theorem 1. *For $sem \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}, st\}$, Algorithm 1 computes $P_{\mathcal{F}}^{sem}(S)$ for both s - and d - prBAFs in time $O(2^{|\mathcal{F} \cdot \mathcal{A}_e| + |\mathcal{F} \cdot \mathcal{E}_e|} \times F(|\mathcal{F}|))$, F is a polynomial function.*

Proof. The fact that Algorithm 1 computes $P_{\mathcal{F}}^{sem}(S)$ follows from the fact that it computes $\sum_{\mathcal{F}^* \in \mathcal{E}(\mathcal{F})} P(\mathcal{F}^*|\mathcal{F}) \times P_{\mathcal{F}^*}^{sem}(S)$. Specifically, Lemma 1 ensures that $P_{\mathcal{F}}^{sem}(S)$ can be computed as $\sum_{\mathcal{F}^* \in \mathcal{E}(\mathcal{F})} P(\mathcal{F}^*|\mathcal{F}) \times P_{\mathcal{F}^*}^{sem}(S)$, where for each $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$. Moreover, for each $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$ Algorithm 1 computes $P_{\mathcal{F}^*}^{sem}(S)$ as specified by Lemma 2.

Finally, it is straightforward to see that computing $P_{\mathcal{F}^*}^{sem}(S)$ as done by Algorithm 1 (i.e., following Lemma 2 and applying the formulas reported in Facts 1 and 2) is feasible in time $O(F(|\mathcal{F}|))$, where F is a polynomial function. Hence, since $|\mathcal{E}(\mathcal{F})| \leq 2^{|\mathcal{F} \cdot \mathcal{A}_e| + |\mathcal{F} \cdot \mathcal{E}_e|}$, it follows that Algorithm 1 runs in time $O(2^{|\mathcal{F} \cdot \mathcal{A}_e| + |\mathcal{F} \cdot \mathcal{E}_e|} \times F(|\mathcal{F}|))$, which completes the proof. \square

4.2 Experimental validation

In this section we report a preliminary experimental assessment of the efficiency of Algorithm 1. To this end we compared running times of Algorithm 1 (denoted as CONTRACT in what follows) with a naive algorithm that computes $P_{\mathcal{F}}^{sem}(S)$ by directly applying Equation 2 of Definition 10 (denoted as NAIVE in what follows).

We perform experiments over 100 prBAFs with a number of arguments ranging over $\{6, 8, 10, 12, 14\}$. Specifically we randomly generate 20 prBAFs for every number of arguments in $\{6, 8, 10, 12, 14\}$. Figure 2 reports the average running times of CONTRACT and NAIVE vs the number of arguments in the prBAFs.

From the experiments it follows that CONTRACT outperforms NAIVE for all the considered number of arguments. However, it is worth noting that even using CONTRACT computing $P_{\mathcal{F}}^{sem}(S)$ requires a large amount of time (the algorithm was halted after 30 minutes) on prBAFs with more than 14 arguments.




Fig. 2: Average running times of CONTRACT and NAIVE (msec) vs number of arguments

5 Conclusions

In this paper we devised an algorithm for computing extensions' probabilities over prBAFs of type IND when the stable, d-admissible, s-admissible or c-admissible semantics are considered. The correctness of the algorithm has been formally proved and its efficiency experimentally validated w.r.t. the naive computation based on the enumeration of possible BAFs. The gain in efficiency of the proposed algorithm is due to the fact that it enumerates contractions rather than possible BAFs and in most cases the number of possible contractions is much smaller than the number of possible BAFs. However, from the experiments it turns out that the algorithm is not able to deal with large prBAFs in reasonable time. Hence, for large prBAFs resorting to estimating extensions' probabilities is reasonable. An interesting research direction for future work is that of applying the approach based on enumerating contractions to improve the efficiency of estimation algorithms.

References

1. Amgoud, L., Vesic, S.: A new approach for preference-based argumentation frameworks. *A. Math. Artif. Intell.* **63**(2), 149–183 (2011)
2. Baroni, P., Giacomin, M.: Semantics of abstract argument systems. In: *Argum. in Artif. Intell.*, pp. 25–44 (2009)
3. Bench-Capon, T.J.M.: Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.* **13**(3), 429–448 (2003)
4. Boella, G., Gabbay, D.M., van der Torre, L.W.N., Villata, S.: Support in abstract argumentation. In: *Proc. of COMMA 2010*. pp. 111–122 (2010)
5. Brewka, G., Woltran, S.: Abstract dialectical frameworks. In: *Proc. of KR (2010)*, <http://aaai.org/ocs/index.php/KR/KR2010/paper/view/1294>
6. Cayrol, C., Lagasque-Schiex, M.C.: On the acceptability of arguments in bipolar argumentation frameworks. In: *Proc. of ECSQARU*. pp. 378–389 (2005)
7. Cayrol, C., Lagasque-Schiex, M.C.: Bipolar abstract argumentation systems. In: *Argum. in Artif. Intell.*, pp. 65–84 (2009)

8. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolarity in argumentation graphs: Towards a better understanding. *Int. J. Approx. Reasoning* **54**(7), 876–899 (2013)
9. Cohen, A., García, A.J., Simari, G.R.: Backing and undercutting in abstract argumentation frameworks. In: *Proc. of FoIKS*. pp. 107–123 (2012)
10. Cohen, A., Gottifredi, S., García, A.J., Simari, G.R.: A survey of different approaches to support in argumentation systems. *Know. Eng. Review* **29**(5), 513–550 (2014)
11. Coste-Marquis, S., Koniec-zny, S., Marquis, P., Ouali, M.: Weighted attacks in argumentation frameworks. In: *Proc. of KR* (2012)
12. Doder, D., Woltran, S.: Probabilistic argumentation frameworks - A logical approach. In: *Proc. of SUM*. pp. 134–147 (2014)
13. Dondio, P.: Toward a computational analysis of probabilistic argumentation frameworks. *Cybernetics and Systems* **45**(3), 254–278 (2014)
14. 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)
15. Dung, P.M., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. *Artif. Intell.* **171**(10-15), 642–674 (2007)
16. Dung, P.M., Thang, P.M.: Towards (probabilistic) argumentation for jury-based dispute resolution. In: *Proc. of COMMA*. pp. 171–182 (2010)
17. Dunne, P.E.: The computational complexity of ideal semantics. *Artif. Intell.* **173**(18) (2009)
18. Dunne, P.E., Hunter, A., McBurney, P., Parsons, S., Wooldridge, M.: Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.* **175**(2) (2011)
19. Dunne, P.E., Wooldridge, M.: Complexity of abstract argumentation. In: *Argum. in Artif. Intell.*, pp. 85–104 (2009)
20. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation. In: *Proc. of IJCAI* (2013)
21. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.* **16**(3), 22 (2015)
22. Fazzinga, B., Flesca, S., Furfaro, F.: Probabilistic bipolar abstract argumentation frameworks: complexity results. In: *Proceedings of IJCAI 2018*. pp. 1803–1809 (2018)
23. Fazzinga, B., Flesca, S., Furfaro, F.: Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence. *Artificial Intelligence* **268**, 1 – 29 (2019). <https://doi.org/https://doi.org/10.1016/j.artint.2018.11.003>
24. Hunter, A.: Some foundations for probabilistic abstract argumentation. In: *Proc. of COMMA*. pp. 117–128 (2012)
25. Hunter, A.: A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reasoning* **54**(1), 47–81 (2013)
26. Hunter, A.: Probabilistic qualification of attack in abstract argumentation. *Int. J. Approx. Reasoning* **55**(2), 607–638 (2014)
27. Hunter, A., Thimm, M.: Probabilistic argumentation with epistemic extensions. In: *Proc. of DARE@ECAI* (2014)
28. Hunter, A., Thimm, M.: Probabilistic argumentation with incomplete information. In: *Proc. of ECAI*. pp. 1033–1034 (2014)
29. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: *Proc. of TAFA* (2011)
30. Martínez, D.C., García, A.J., Simari, G.R.: On acceptability in abstract argumentation frameworks with an extended defeat relation. In: *Proc. of COMMA*. pp. 273–278 (2006)
31. Modgil, S.: Reasoning about preferences in argumentation frameworks. *Artif. Intell.* **173**(9-10), 901–934 (2009)
32. Nouioua, F., Risch, V.: Bipolar argumentation frameworks with specialized supports. In: *Proc. of ICTAI*. pp. 215–218 (2010)
33. Oren, N., Norman, T.J.: Semantics for evidence-based argumentation. In: *Proc. of COMMA*. pp. 276–284 (2008)

34. Polberg, S., Hunter, A.: Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *Int. J. Approx. Reasoning* **93**, 487–543 (2018)
35. Proietti, C.: Polarization and bipolar probabilistic argumentation frameworks. In: *Proc. of Work. AI³*. pp. 22–27 (2017)
36. Rienstra, T.: Towards a probabilistic dung-style argumentation system. In: *Proc. of AT*. pp. 138–152 (2012)
37. Thimm, M.: A probabilistic semantics for abstract argumentation. In: *Proc. of ECAI*. pp. 750–755 (2012)
38. Verheij, B.: The toulmin argument model in artificial intelligence. In: *Argum. in Artif. Intell.*, pp. 219–238 (2009)