Conditional Probabilistic Bipolar Argumentation Framework: Explanations, Complexity and Approximation

Gianvincenzo Alfano, Sergio Greco, Domenico Mandaglio, Francesco Parisi, Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering,
University of Calabria, ITALY

{g.alfano, greco, d.mandaglio, fparisi, i.trubitsyna}@dimes.unical.it

Università della Calabria
DIPARTIMENTO DI INGEGNERIA INFORMATICA, MODELLISTICA, ELETTRONICA E SISTEMISTICA
DIMES

FAIR Future Artificial Intelligence Research

SERICS

Bipolar Argumentation Framework (BAF)

- A Bipolar Argumentation Framework (BAF) is a triple $\langle A, R, S\rangle$ where:
  - A is a set of arguments
  - $R \subseteq A \times A$ is a finite set of attacks
  - $S \subseteq A \times A$ is a finite set of supports
- It allows representing dialogues, making decisions, and handling inconsistency
- Can be viewed as a direct graph: nodes are arguments, edges are attacks/supports

Semantics for BAFs

- An argumentation semantics σ specifies the criteria for identifying "reasonable" sets of arguments, called *extensions*
- A *complete extension* is an admissible set that contains all the arguments that it defends
- A complete extension E is said to be:
  - *preferred* (pr) iff it is maximal (w.r.t. $\subseteq$)
  - *stable* (st) iff it attacks all arguments in $A \setminus E$
  - *grounded* (gr) iff it is minimal (w.r.t. $\subseteq$)

Probabilistic BAF

- Extends BAF with a non-zero marginal probability for every BAF element (i.e. argument/attack/support)
- $P(x)$ represents quantified uncertainty about the occurrence of BAF element x
- Semantics defined through *possible worlds*: BAFs obtained by removing consistent subsets of the probabilistic elements
- Can't encode $i)$ cyclic supports (though the semantics of cyclic BAF has been investigated), and $ii)$ conditional probabilities

Conditional Probabilistic BAF

Contribution: We propose the *Conditional Probabilistic Bipolar Argumentation Framework* (CPBAF), that is a quadruple $\Delta = \langle A, R, S, P\rangle$ where $\langle A, R, S\rangle$ is a (cyclic) *BAF*, and P is a total conditional probability function assigning a non-zero probability value $P(a|C_a)$ to every element $a \in A \cup R \cup S$, where C_a is a conditional event consisting of a propositional logic formula whose atoms are taken from $A \cup R \cup S$.

Example 1. Consider the CPBAF $\langle A = \{\mathtt{a},\mathtt{b},\mathtt{c},\mathtt{d},\mathtt{e}\}, R = \{(\mathtt{a},\mathtt{b}),(\mathtt{b},\mathtt{e}),(\mathtt{c},\mathtt{d}),(\mathtt{d},\mathtt{c})\}, S = \{(\mathtt{b},\mathtt{c}),(\mathtt{e},\mathtt{a})\}, P\rangle$ with possible worlds $w_1 \dots w_5$.

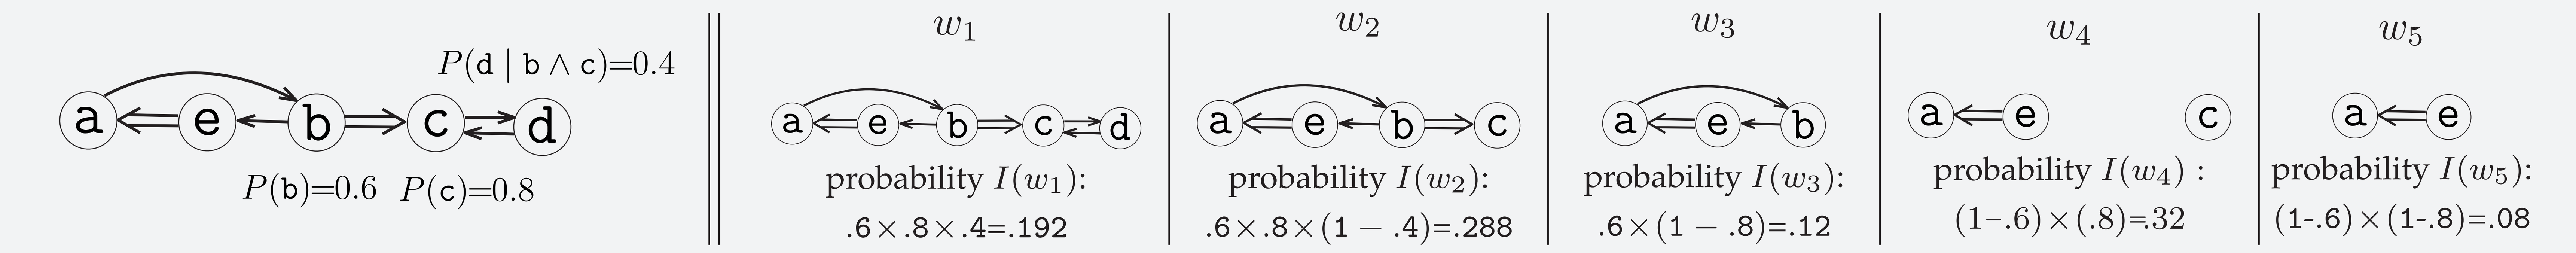

Contribution: We define the Probabilistic Acceptance PrA$[\sigma]$ problem in CPBAF. Given a CPBAF Δ and an argument g, the problem asks for the probability that g is accepted in Δ, by means of some fixed Prob. Distr. Func. (PDF $Pr(\cdot, w, \sigma)$) over the σ-extensions of the possible worlds of Δ.

$$PrA^{\sigma}_{\Delta}(g) = \sum_{\substack{w \in pw(\Delta) \wedge \\ E \in \sigma(w) \wedge g \in E}} I(w) \cdot Pr(E, w, \sigma)$$

$$I(w) = \prod_{t \in T' \wedge w \models C_t} P(t \mid C_t) \times \prod_{t \in T \setminus T' \wedge w \models C_t} (1 - P(t \mid C_t)),$$

where T = A $\cup R \cup S$ and $T' = A' \cup R' \cup S'$.

Theorem 4.
For $\sigma \in \{\mathtt{gr},\mathtt{co},\mathtt{st},\mathtt{pr}\}$, PrA$[\sigma]$ is FP$^{\#P}$-hard, even for acyclic CPBAFs and for any chosen PDF.

Explanations for BAF

Contribution: We introduce our notion of explanation for a BAF Λ, and exploit it to provide a PDF over the σ-extensions of Λ. This leads to an instantiation of PrA$[\sigma]$, dubbed PrEA$[\sigma]$. Our explanation strategy satisfies properties defined in the literature, such as σ-basic, σ-existence, Conflict-freeness, Relevance, Disjointness, Sequence-Minimality, and Inclusion.

The computation of an explanation can be carried out considering one SCC at a time following the topological order of the argumentation graph, and for each SCC by alternating 'deterministic computations' (using the grounded semantics) and the choice of an argument to be accepted.

Example 2. Consider the BAF Λ, where $\langle \mathtt{j}, \mathtt{a}\rangle$ is an explanation for the complete extension $E = \{\mathtt{a},\mathtt{d},\mathtt{e},\mathtt{i},\mathtt{j}\}$.

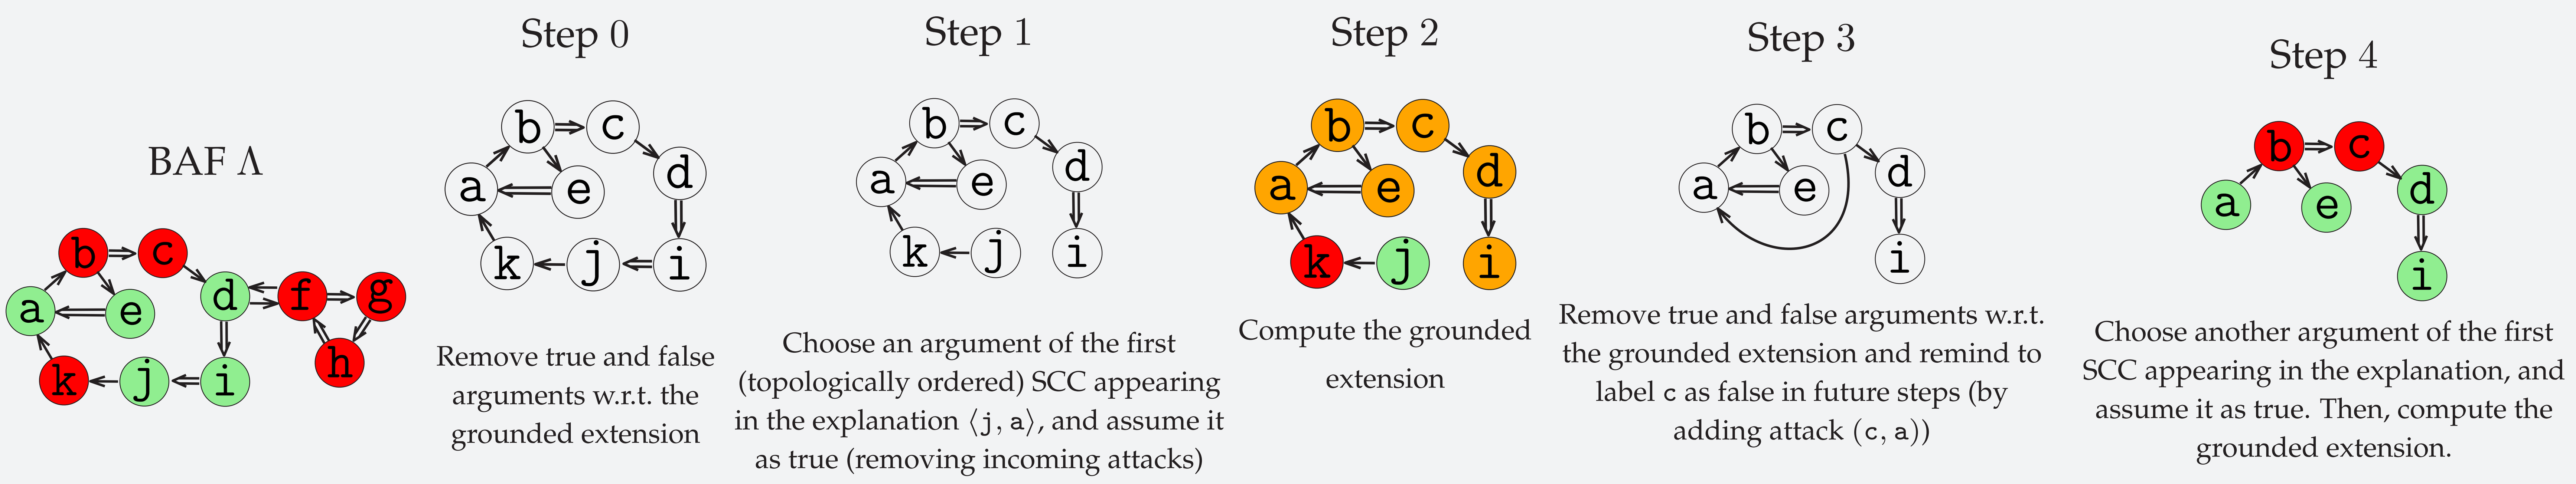

- Since extensions may have multiple explanations of different lengths, it is reasonable to assume that some explanations are preferred to others
- To define probabilities of explanations, we exploit the concepts of probabilistic *trie* (standard prefix tree data structure)
- As extensions do not share explanations, the probability of explanations is transferred to the extensions and acceptance of arguments
- For the CPBAF of Example 1, $PrEA_{\Delta}(g) = 1/3$ (resp., $1/3$, $1/9$, $4/9$, and $1/3$), with $g = \mathtt{a}$ (resp., b, c, d, and e)

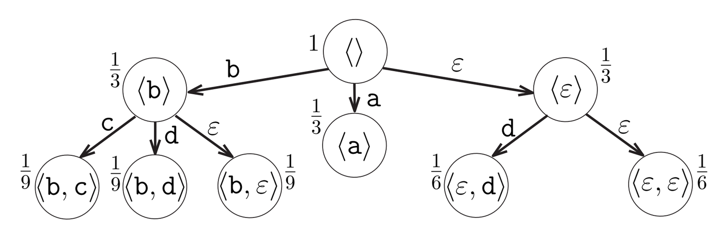

Approximations

Contribution: To deal with the intractability of PrA$[\sigma]$ and PrEA$[\sigma]$, whose complexity is shown to be FP$^{\#P}$-hard, we propose an *additive approximation algorithm* for PrEA$[\sigma]$ for CPBAF without cycles with an odd number of attacks and semantics $\sigma \in \{\mathtt{gr},\mathtt{co},\mathtt{st},\mathtt{pr}\}$.

Theorem 4. Unless NP $\subseteq BPP$, there is no FPRAS for PrA$[\sigma]$ with $\sigma \in \{\mathtt{gr},\mathtt{co},\mathtt{st},\mathtt{pr}\}$, even for acyclic CPBAFs.

Theorem 5. Unless NP $\subseteq BPP$, there is no FPARAS for PrA[st] and PrA[pr], for any chosen PDF.

Theorem 6. PrEA$[\sigma]$ has an FPARAS if i) σ=gr, or ii) $\sigma \in \{\mathtt{gr},\mathtt{co},\mathtt{st},\mathtt{pr}\}$ and the input CPBAF is odd-cycle-free.

Theorem 7. Whenever i) $\sigma =$, or ii) $\sigma \in \{\mathtt{gr},\mathtt{co},\mathtt{st},\mathtt{pr}\}$ and the input BAF Λ is odd-cycle-free, then: Algorithm 2 runs in polynomial time and, for each $E \in \sigma(\Lambda)$, it outputs E with probability $Pr(E, \Lambda, \sigma)$.

Algorithm 1: Apx

Input: A CPBAF Δ, a semantics σ, a goal $g \in A$, error parameter $\epsilon > 0$, and uncertainty parameter $0 < \delta < 1$.
Output: a random number p s.t. $PrEA^{\sigma}_{\Delta}(g) \in [p-\epsilon, p+\epsilon]$ with probability $1 - \delta$.

```
n = ⌈1/(2ε²) × ln(2/δ)⌉; κ = 0;
for i ∈ {1, ..., n} do
    Choose Λ=⟨A, R, S⟩ in pw(Δ) with probability I(Λ);
    Choose X ∈ ξ^σ(Λ) with probability π(X);
    if g ∈ gr(⟨A, R \ (A × X̃), S \ (A × X̃)⟩ then
        κ = κ + 1;
return κ/n;
```

Algorithm 2:

Input: A BAF $\Lambda = \langle A, R, S\rangle$ and a semantics σ.
Output: An explanation for a σ-extension.

```
Let X = ⟨⟩;  F = U = ∅;  Λ = Λ_{↑gr(Λ)*};
while Λ ≠ ⟨∅, ∅, ∅⟩ do
    Let A′ be the first SCC of Λ;
    Let C = {a ∈ A′ | a ∉ (F ∪ U) ∧ (∃x ∈ A′.(a, x) ∈ R) ∧ ((F ≠ ∅) (gr(Λ_a)^+ ∩ F ≠ ∅))};
    if σ = gr then C = {ε};
    if σ = co ∧ F = ∅ then C = C ∪ {ε};
    Select a ∈ C with probability 1/|C| and append it to X;
    if a = ε then
        U = U ∪ {x | (y, x) ∈ ((R∪S) ∩ (A′ × (A \ A′))}; Λ = Λ_{↑A′};
    else F = F_a; and Λ = Λ̂_a;
return X
```