IJCAI 2026 · MAIN TRACK

When Incompleteness Does Not Matter

The case of Incomplete Abstract Argumentation Frameworks, under the possible perspective

Bettina Fazzinga[1] · Sergio Flesca[2] · Filippo Furfaro[2]

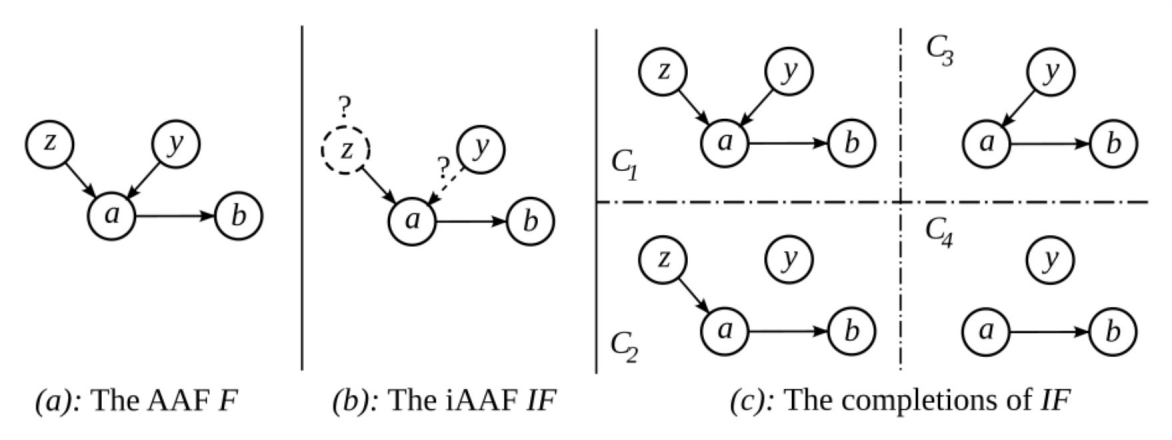
[1] DICES – University of Calabria · [2] DIMES – University of Calabria

A motivating scenario

Lawyer Ann models a trial as an argumentation framework: witness and expert statements are arguments, perceived conflicts between them are attacks.

She is not sure whether witness **z** will testify, or whether the jury will perceive the attack between **y** and **a**.

So she marks z and (y,a) as uncertain — an iAAF with four possible completions.



The plain AAF (a), the iAAF (b, dashed = uncertain), and its four completions (c)

Reasoning under the *possible* perspective, Ann learns that argument **a** — dismissed by the plain AAF — is possibly accepted after all. That extra insight is exactly what an iAAF buys her.

Background: Abstract Argumentation

The framework

An AAF is a pair F = ⟨A, D⟩: arguments A and attacks D between them (Dung, 1995).

A semantics σ selects sets of arguments — **extensions** — that can coherently stand together.

The Dungean semantics

ad admissible · **co** complete · **gr** grounded · **pr** preferred · **st** stable

An argument is **Cr-** (credulously) accepted if it is in some extension, **Sk-** (skeptically) accepted if it is in every extension.

AAFs are flexible and well understood — but they are also **certain**: every argument and attack either is, or is not, in the graph.

Incomplete AAFs (iAAFs)

An iAAF is a tuple **IF = ⟨A, A?, D, D?⟩**: certain arguments A and attacks D, plus uncertain arguments A? and attacks D? that may or may not occur.

arg-iAAF: only arguments are uncertain (D? = ∅). · att-iAAF: only attacks are uncertain (A? = ∅).

Completions

A completion is a plain AAF ⟨A′, D′⟩ obtained by resolving every uncertain component: keep it, or drop it.

WHY THIS MATTERS

A small iAAF compactly packs an **exponential** number of completions — up to 2^|A?| × 2^|D?| alternative scenarios.

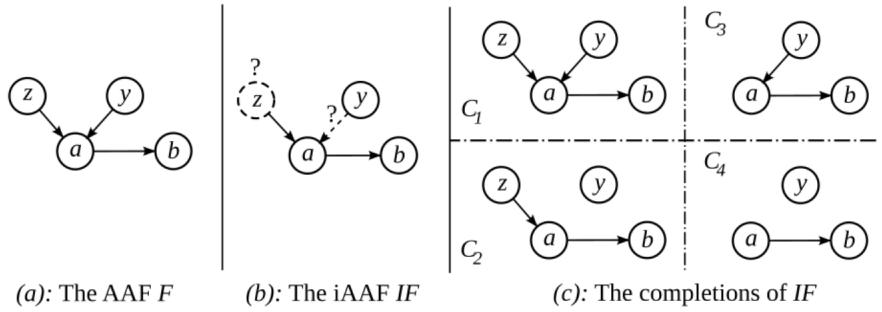
Possible vs. necessary reasoning

Possible perspective

S is an extension, a is accepted, if this holds in **at least one** completion.

Necessary perspective

S is an extension, a is accepted, if this holds in **every** completion. (Left to future work.)



(a): The AAF *F* *(b):* The iAAF *IF* *(c):* The completions of *IF*

Ann's iAAF IF (b) and its four completions (c)

BACK TO ANN'S IAAF, UNDER ADMISSIBLE SEMANTICS

Possible extensions: {Ø, {y}, {z}, {y,b}, {z,b}, {a}, {y,a}}

Necessary extensions: only {Ø, {y}}

Every argument y, z, a, b is *possibly* credulously accepted — but *none* is skeptically accepted, possibly or necessarily.

This paper studies the possible perspective throughout.

THE OPEN QUESTION

Does adding qualitative uncertainty to AAFs genuinely increase expressive power?

Given an iAAF IF, is there always an AAF F such that reasoning over the possible extensions of IF can be carried out by classical, extension-based reasoning over F?

Two notions of correspondence

EQUIVALENCE — IF ≡ F

The extensions of F are **exactly** the possible extensions of IF.

Any auxiliary argument introduced to simulate uncertainty is purely a support device — it never appears in an extension.

PROJECTION-EQUIVALENCE — IF ≡Π F

The possible extensions of IF are obtained by **projecting away** the auxiliary arguments of F.

Auxiliary arguments may end up inside F's extensions, but are simply ignored when projecting back.

Equivalence ⇒ projection-equivalence, not vice versa. Whenever a construction exists, we insist on **linear size** in |IF| — not just "some" AAF, however large.

The picture, semantics by semantics

Semantics σ	Equivalent AAF?	Projection-equivalent AAF?
admissible (ad)	✓ Always — linear time	✓ Yes (implied)
stable (st)	✗ Not always	✓ Always — linear time
preferred (pr)	✗ Not always	✓ Always — linear time
complete (co)	✗ Not always	✗ Not always
grounded (gr)	✗ Not always	✗ Not always

Every positive result comes with a constructive, linear-time algorithm producing an AAF of linear size.

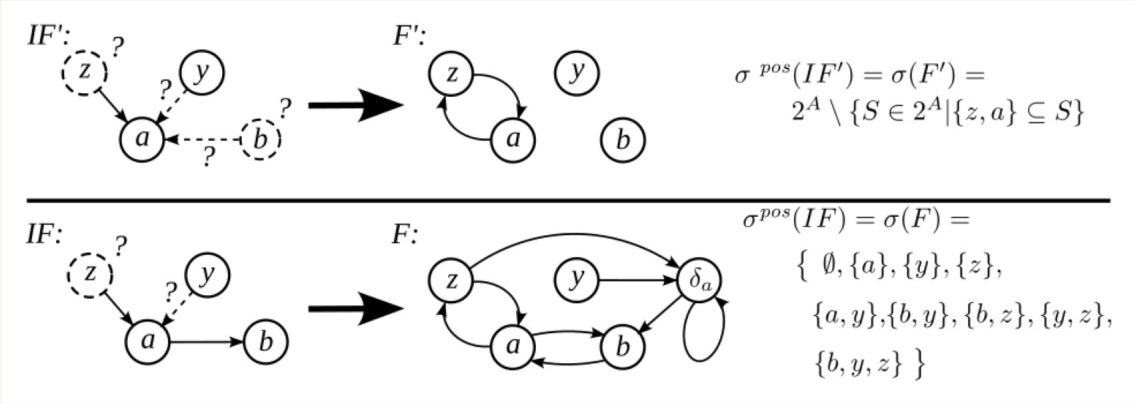
Admissible: uncertainty costs nothing

1. An uncertain argument attacking nothing with certain attacks: just make it certain — admissible semantics never forces its inclusion.

2. An uncertain argument attacking some argument with a certain attacks: make it certain, add counterattacks back onto it.

3. An uncertain attack into a "sink" argument: harmless — simply drop it.

4. An uncertain attack (y,a) where a attacks further: replace it with a self-attacking dummy δ_a reproducing the same defenses.



Algorithm 1 applied to two iAAFs (left → right), with the resulting extensions

Theorem 1. Algorithm 1 runs in linear time and returns an AAF F with IF $\equiv^{ad}$ F.

Beyond admissible: equivalence breaks

THEOREM 2

Under **st, co, gr, pr**, there exist iAAFs — even a 3-argument arg-iAAF or att-iAAF — with **no** equivalent AAF, of any size whatsoever.

Why co fails

A={a,b,c}, D={(b,c)}, D?={(a,b)}: possible extensions {a,b} and {a,c} intersect to {a} — not itself a possible extension. Yet complete extensions of any AAF are always closed under intersection.

Why st, pr, gr fail

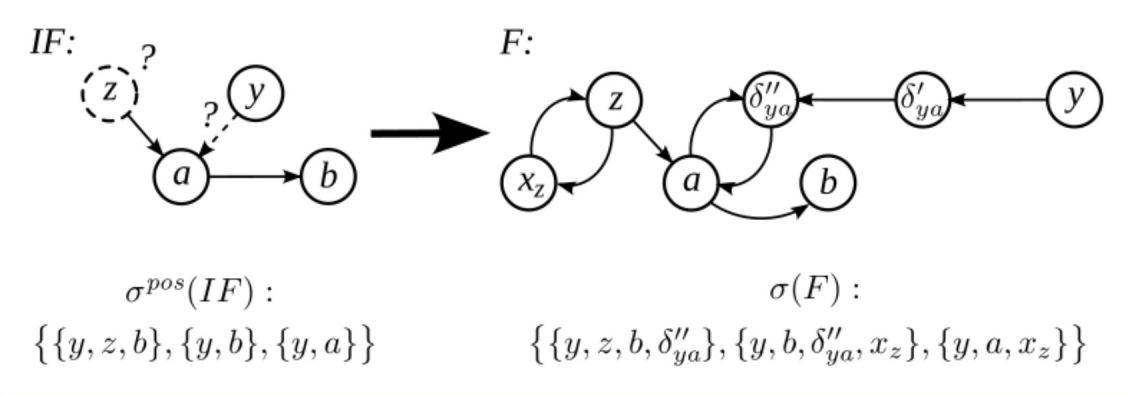
A={a}, A?={b}, D=∅: possible extensions {a} ⊂ {a,b} are nested — impossible for stable/preferred extensions of any AAF, and grounded always yields exactly one extension.

Stable: encode the choice, then project

In a stable extension S, every argument is either in S, or attacked by something in S.

So: make each uncertain argument a certain, in mutual attack with a fresh x_a. Extensions without a now contain x_a instead.

Each uncertain attack (a,b) becomes a small certain gadget (δ′, δ″) that lets both "with" and "without (a,b)" appear as genuine stable extensions.



Algorithm 2: three possible extensions of IF ↔ three stable extensions of F

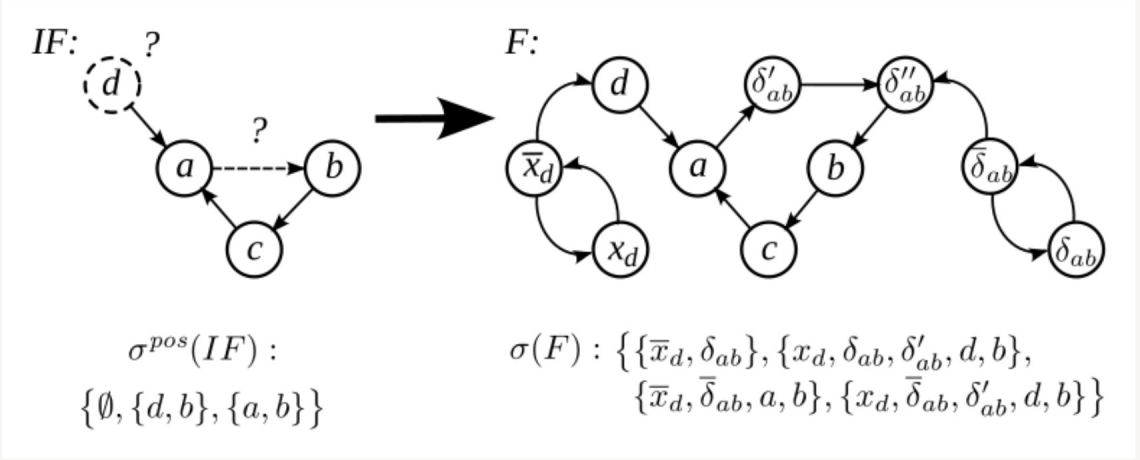
Linear time, linear size — and here a genuine one-to-one bijection between possible extensions of IF and extensions of F.

Preferred: simulate every completion

For each uncertain argument a: add a mutually-attacking pair x_a, $\bar{x}_a$, with $\bar{x}_a$ attacking a — a preferred extension now "chooses" x_a (a in) or $\bar{x}_a$ (a out).

Each uncertain attack (a,b) gets an analogous chosen/not-chosen gadget.

F now literally encodes every completion of IF.



Algorithm 3: possible extension {d,b} maps to two extensions of F

Still linear time, linear size — but no longer a bijection: one possible extension of IF may correspond to **many** extensions of F.

Complete and grounded: no way out

THEOREM 3

Under **co** and **gr**, there exist iAAFs with **no** projection-equivalent AAF either — however large F is allowed to be.

Even projecting away auxiliary arguments cannot repair the intersection-closure property that complete extensions need, or give more than one grounded extension.

Takeaway: under co and gr, qualitative uncertainty is genuine extra expressive power — it cannot be compiled away.

What the transformations guarantee

LINEAR SIZE

|F| is linear in |IF| for all three algorithms (ad, st, pr) — despite an exponential number of completions.

PRESERVATION

Every original argument and every certain attack of IF still occurs in F.

ARGUMENT-INVARIANCE

Admissible only: when uncertainty involves just arguments, or attacks into sinks, F adds **no new arguments** at all — F stays human-readable.

Reusing standard AAF machinery

ad, st	A genuine bijection lets verification, enumeration *and* counting of possible extensions reduce directly to off-the-shelf AAF solvers on F.
pr	Verification and enumeration still possibly done over F — but **not counting**: one possible extension can correspond to exponentially many extensions of F.
psa, nca	These need to see the individual completions, not just the possible extensions — so they are **not** solved by this machinery.

Credulous acceptance under the possible perspective always transfers, for ad, st and pr alike.

Related work

Alfano et al. (2023) reduced iAAFs to simpler arg-/att-iAAFs preserving possible extensions under st, gr, co, pr. Our st/pr results go further: uncertainty is removed *entirely*. Our co/gr negative result strengthens theirs: full removal is provably impossible, not just harder.

Complexity results (**Fazzinga et al.**) show verification is in P for ad, st, gr, co — same as AAFs — and Σ_2^p-complete for pr, one level higher. Matching complexity does not guarantee a compact transformation exists (co, gr), and mismatched complexity does not rule one out (pr).

Signature results (**Dunne et al.**) show $\Sigma_a d$ is closed under union, so *some* equivalent AAF always exists under admissible — but say nothing about its size or how to build it. We give a linear-time construction.

Conclusions & future work

ad: no extra power at all — an equivalent AAF always exists, in linear time.

st, pr: no extra power once auxiliary arguments are allowed — a projection-equivalent AAF always exists, still linear.

co, gr: genuine extra power — uncertainty cannot be compiled away, at any size.

When correspondence holds, classical AAF solvers can be reused directly — no bespoke iAAF machinery needed.

Future: the necessary perspective; more compact/interpretable heuristics; experiments against native iAAF solvers; implications for explainable AI.

Thank you

Questions and discussion welcome.

Bettina Fazzinga · Sergio Flesca · Filippo Furfaro

bettina.fazzinga, sergio.flesca, filippo.furfaro @ unical.it