When Incompleteness Does Not Matter: The Case of Incomplete Abstract Argumentation Framework (Under the Possible Perspective)

Bettina Fazzinga[1] , **Sergio Flesca**[2] , **Filippo Furfaro**[2*]
[1]DICES - University of Calabria
[1]DIMES - University of Calabria
{bettina.fazzinga, sergio.flesca, filippo.furfaro}@unical.it

Abstract

Incomplete Abstract Argumentation Frameworks (iAAFs) extend Abstract Argumentation Frameworks (AAFs) by allowing arguments and attacks to be specified as uncertain, enabling a compact representation of alternative argumentation scenarios. Despite extensive work on reasoning with iAAFs, their expressive power relative to standard AAFs is still unclear. We address this question by studying whether the reasoning based on possible-extensions in iAAFs can be reduced to classical reasoning on AAFs. Our analysis relates iAAFs and AAFs via two comparison notions: *equivalence*, where an AAF yields as extensions exactly the possible extensions of an iAAF, and *projection-equivalence*, where such extensions are obtained by projecting away auxiliary arguments. We characterize the semantics under which these relationships hold, and, in these cases, provide constructive transformations from iAAFs to AAFs, thus also enabling classical argumentation tools to be applied to qualitative-uncertainty reasoning.

1 Introduction

Abstract Argumentation Frameworks (AAF) [Dung, 1995] are a well-established AI paradigm, valued for their flexibility in modeling diverse problems—from dispute management to fields like process mining [Fazzinga *et al.*, 2022]. This success has driven research into AAF extensions designed to broaden their scope, particularly those capable of encoding uncertainty. In this regard, *probabilistic AAFs* have been introduced to enrich AAFs with quantitative information about the likelihood of arguments and attacks [Li *et al.*, 2011; Fazzinga *et al.*, 2015; Fazzinga *et al.*, 2019]), thereby representing a probability distribution over the candidate frameworks. While probabilistic AAFs encode the uncertainty quantitatively, *incomplete AAFs* (iAAFs [Baumeister *et al.*, 2018; Fazzinga *et al.*, 2020; Fazzinga *et al.*, 2026]) extend AAFs by introducing *qualitative* uncertainty: arguments or attacks may be marked as *uncertain*, meaning they may or may not occur, without committing to probability values. This simple mechanism makes iAAFs particularly intuitive and user-friendly: an analyst can compactly describe a range of possible argumentation scenarios by labeling some components as potentially absent. As a consequence, several works have examined different aspects of iAAF reasoning, such as the complexity of acceptance and extension verification, algorithmic procedures for evaluating the acceptance status, and generalizations incorporating supports or correlations between uncertain components [Baumeister *et al.*, 2021; Fazzinga *et al.*, 2020; Fazzinga *et al.*, 2021a; Fazzinga *et al.*, 2021b; Fazzinga *et al.*, 2023; Fazzinga *et al.*, 2026; Mailly, 2024].

Example 1. *Lawyer Ann is reasoning about an upcoming trial. To anticipate how the jury might evaluate the evidence, Ann mentally adopts the jury's perspective and constructs an abstract argumentation model of the case. Witness and expert statements are represented as arguments, while the conflicts between the statements (as supposedly perceived by the jurors) as attacks. The resulting argumentation graph is in Figure 1*(a)*. In the attempt to foresee which lines of reasoning the jury could find persuasive, Ann looks into the preferred extensions of the so-obtained argumentation framework, and realizes that the unique preferred extension is* $S_1 = \{y, z, b\}$.

Now, Ann is not sure that the witness that is supposed to claim the argument z *will show up at the trial, and that the jury will perceive the attack between arguments* y *and* a *(since the jurors may not have the technical skills required to perceive this attack). Hence, Ann decides to "upgrade" to the iAAF paradigm, and specifies that argument* z *and attack* (y, a) *are uncertain (see Figure 1*(b)*). This means that the four alternative scenarios for the argumentation graph in Figure 1*(c)*, called* completions*, must be now considered. Correspondingly, the possible and the necessary perspectives*

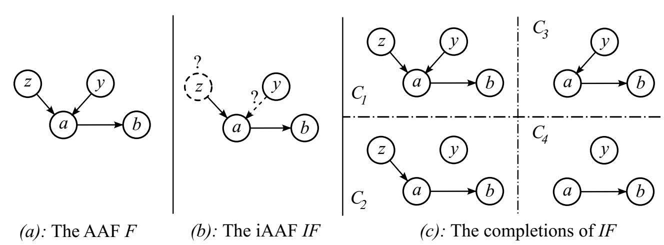

Figure 1: The AAF F and the iAAF IF of Example 1

[*]Contact Author

for the reasoning can be adopted: under the possible (resp., necessary) perspective, a set S is an extension and an argument a is accepted if and only if S is an extension and a accepted in at least one (resp., in every) completion. Thus, by reasoning under the possible perspective, Ann gets the information that, as the uncertainty may be resolved in different ways, not only S_1 should be considered as a preferred extension, but also $S_2 = \{y, b\}$ and $S_3 = \{y, a\}$, that are possible extensions (as they are extensions in C_3 and C_4, respectively). In turn, not only the arguments y, z, b should be considered as accepted, since also a turns out to be possibly accepted (as it is accepted in C_4). So, Ann is now aware that there is some chance that the jury will perceive a as a justified argument, and, if she believes that a is contrary to her client's interests and if she does not want to run the risk of this situation occurring, she may decide to invest in presenting additional arguments that could weaken a. □

Despite the considerable attention devoted to iAAFs, a fundamental question remains open: *does adding qualitative uncertainty to AAFs genuinely increase expressive power*? More precisely, given an iAAF IF, it is unknown whether reasoning based on its *possible* or its *necessary* extensions can be accomplished by reasoning over the extensions of an AAF "equivalent" to IF.

In this paper, we address this gap. Specifically, we focus on the possible perspective of the reasoning and study whether, given an iAAF IF, there exists an AAF F such that: (i) the set of its extensions coincides with the set of possible extensions of IF (*equivalence*); or (ii) the possible extensions of IF can be recovered from the extensions of F by projecting away auxiliary arguments introduced to simulate uncertainty (*projection-equivalence*).

Our results reveal a nuanced picture. Under the admissible semantics, every iAAF admits an equivalent AAF. Under the stable and the preferred semantics, every iAAF admits a projection-equivalent AAF. For other classical semantics (i.e. grounded and complete), we show that neither equivalence nor projection-equivalence necessarily holds. Besides clarifying the expressive relationship between iAAFs and AAFs, our analysis has practical implications: when equivalence or projection-equivalence is achievable, we provide explicit algorithms translating an iAAF into an equivalent (or projection-equivalent) AAF *whose size is linear in the size of the iAAF*. This is in some ways surprising: the completions of an iAAF are exponential in the size of the iAAF, so it is not obvious that traditional reasoning on an AAF of comparable size can account for an exponential number of scenarios. Given this, we show in which cases the so-obtained AAFs can be used as the input of standard argumentation reasoners (working over AAFs) to solve the verification, extension counting or enumeration, and acceptance problems over iAAFs, without the need for an ad-hoc machinery.

Beyond its immediate contribution to abstract argumentation theory, this work has broader implications for several areas of Artificial Intelligence concerned with reasoning under uncertainty. By addressing conditions under which qualitative uncertainty can be compiled away, our results contribute to a deeper understanding of the expressiveness of formalisms with incomplete knowledge: this insight is relevant to knowledge representation and non-monotonic reasoning, where similar questions arise concerning the expressive role of language extensions designed to model incomplete information [Governatori *et al.*, 2004]. Moreover, the ability to reason about incomplete argumentative structures while ultimately relying on classical frameworks can inform the design of reasoning mechanisms in multi-agent systems, where agents often operate with partial views of others' beliefs or interactions [Herzig and Yuste Ginel, 2021].

2 Preliminaries

We assume that the reader is familiar with *Abstract Argumentation Frameworks* (AAFs) and, in particular, with the notions of *extension* and of *credulously* (*Cr*-) and *skeptically* (*Sk*-) accepted argument, under the so-called *Dungean* semantics: *admissible* (ad), *stable* (st), *grounded* (gr), *complete* (co), *preferred* (pr). AAFs will be denoted as pairs of the form $\langle A, D\rangle$, where A and D are the sets of arguments and attacks, respectively. Given an AAF F under a semantics σ, we will denote the set of its extensions as $\sigma(F)$, and the sets of its credulously and skeptically arguments as $\textsc{ca}^\sigma(F)$ and $\textsc{sa}^\sigma(F)$, respectively. Given this, we here review *incomplete Abstract Argumentation Frameworks* (iAAFs) [Baumeister *et al.*, 2018], that allow the uncertainty affecting the presence of arguments and attacks to be qualitatively modeled.

Definition 1. *An* incomplete AAF *(iAAF) is a tuple $IF = \langle A, A^?, D, D^?\rangle$, where A and $A^?$ are disjoint sets of arguments, and D and $D^?$ disjoint sets of attacks between arguments in $A \cup A^?$. The arguments and attacks in A and D (resp., $A^?$ and $D^?$) are said to be* certain *(resp.,* uncertain*), since they are (resp., are not) guaranteed to occur. If $D^? = \emptyset$ (resp., $A^? = \emptyset$), IF is called an* argument-incomplete AAF*, or arg-iAAF (resp., an* attack-incomplete AAF*, or att-iAAF).*

iAAFs compactly represent the alternative scenarios for the argumentation graph, i.e. the possible combinations of arguments and attacks that can occur according to what is certain and uncertain. Each scenario is an AAF called *completion*.

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

The possible and necessary perspectives are natural ways to account for the presence of multiple completions when adapting the notions of *extension* and *acceptance* to iAAFs:

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

Under a semantics σ, we denote the sets of possible and necessary extensions of an iAAF IF as $\sigma^{\text{pos}}(IF)$ and $\sigma^{\text{nec}}(F)$, the sets of its credulously and skeptically accepted arguments as $\textsc{pca}^\sigma(IF)$ and $\textsc{psa}^\sigma(IF)$, under the possible perspective, and as $\textsc{nca}^\sigma(IF)$ and $\textsc{nsa}^\sigma(IF)$, under the necessary perspective, respectively.

Example 2. *Under $\sigma = ad$, the possible extensions of the iAAF IF in Figure 1(b) are $\sigma^{pos}(IF) = \{\, \emptyset, \{y\}, \{z\}, \{y, b\}, \{z, b\}, \{a\}, \{y, a\}\}$ (in particular, $\{a\}$, $\{y, a\}$ are extensions of C_4). Among these, only $\emptyset$ and $\{y\}$ are necessary extensions. Moreover, $\text{PCA}^{ad} = \{y, z, a, b\}$, i.e. every argument is possibly Cr-accepted, since it is Cr-accepted in at least one completion. Indeed, $\text{PSA}^{ad} = \text{NSA}^{ad} = \emptyset$, i.e. no argument is Sk-accepted (neither possibly nor necessarily), since $\emptyset$ is an extension of every completion.*

3 Problem Statement

Inspired by the motivations outlined in the introduction, we seek to provide insights into whether there is a difference in expressiveness between the iAAF and AAF paradigms: more specifically, whether it is feasible to simulate the reasoning over iAAFs based on the adaption of the notion of extension (reviewed in Section 2 - Definition 3) via the traditional extension-based reasoning over AAFs. In particular, we focus on the possible perspective of the reasoning, and defer the study of the necessary perspective to future work.

To this end, we introduce two notions of equivalence between iAAFs and AAF (in doing this, we use the *projection* operator Π defined as follows: given two sets of arguments A, S, and a set of sets $\mathcal{S}$, $\Pi_A(S)$ is the set $S \cap A$, and $\Pi_A(\mathcal{S})$ the set $\{S' | \exists X \in \mathcal{S} \text{ s.t. } \Pi_A(X) = S'\}$).

Definition 4. *Let $IF = \langle A, A^?, D, D^? \rangle$ be an iAAF, $F = \langle A', D' \rangle$ an AAF, and σ a semantics. IF and F are equivalent, written as $IF \equiv^\sigma F$ (resp., projection-equivalent, written as $IF \equiv^\sigma_\Pi F$) if $\sigma^{pos}(IF) = \sigma(F)$ (resp., $\sigma^{pos}(IF) = \Pi_{A \cup A^?}\sigma(F)$).*

Obviously, $IF \equiv^\sigma F$ implies $IF \equiv^\sigma_\Pi F$. Given this, the problem addressed in the rest of the paper is:

"*Let σ be a semantics. Does there always exist, for any iAAF IF, an AAF F such that $IF \equiv^\sigma F$ or, at least, $IF \equiv^\sigma_\Pi F$?*"

Indeed, in posing this question we impose no restriction on how equivalent or projection-equivalent AAFs may be constructed. The only obvious constraint, stemming from the notion of equivalence, is that such AAFs must contain all arguments occurring in the possible extensions of the iAAF. That said, the difference between equivalence and projection-equivalence is in the role of any new arguments introduced to define parts of the AAF that simulate uncertainty. In the case of equivalence, these new arguments serve only as auxiliary devices to support the reasoning process and do not belong to the extensions. In contrast, under projection-equivalence, such arguments may end up being integral parts of the extensions, although they can be ignored when projecting the extensions to obtain the possible extensions. Apart from this distinction, the problem formulation leaves complete freedom to add or remove arguments and attacks, allowing the construction of equivalent AAFs of arbitrary size. In answering this question, however, we will be precise: when we claim a positive result, we show that an equivalent AAF of comparable size -namely, linear size- can be constructed. When we claim a negative result, we mean that no such construction exists, even imposing no bound on the size of the AAF.

4 When Do IAAFs Admit Equivalent or Projection-Equivalent AAFs?

We address the fundamental research question introduced so far starting from the case of *admissible* semantics. In this case, we introduce Algorithm 1, that translates any input iAAF IF into an equivalent AAF F. The underlying rationale is summarized by the following intuitions, that suggest how the uncertainty in an iAAF can be "reduced", by reformulating uncertain arguments/attacks as (groups of) certain arguments and attacks:

$-\mathcal{I}_1$: the semantics of an iAAF (i.e. the set of possible extensions) does not change if every uncertain argument a that is source of no certain attack is converted into a certain argument. The reason is twofold: 1) the fact that all attacks stemming from a (if any) are uncertain suffices to model the case where no argument is attacked by a or defended by a; 2) even if a is made certain, the possibility that a is not included in some extensions is still guaranteed by the admissible semantics, that does not force the arguments that are acceptable w.r.t an extension to be included in the extension itself.

$-\mathcal{I}_2$: also the case where an uncertain argument a is the source of certain attacks can be simulated by making a certain, but now also a counterattack (b, a) must be introduced for each certain attack (a, b) stemming from a. This allows any argument b attacked by a to be in a possible extension not containing a, if this is not forbidden by the relationship of b with the arguments other than a; such a possibility, when a was uncertain, was granted in the completions without a.

$-\mathcal{I}_3$: uncertain attacks towards arguments attacking no argument have no impact on the set of possible extensions, so they can be discarded. The reason is that removing these attacks does not discard any defense in any scenario (so it does not deprive any argument from a chance of being accepted in some scenario), and preserves the possibility that the target arguments of these attacks belong to some possible extension.

$-\mathcal{I}_4$: uncertain attacks towards arguments attacking some argument, on the contrary, have an impact on the set of possible extensions. However, in this case, the effects of the uncertain attacks (y, a) can be simulated by creating a self-contradicting argument δ_a and performing the following actions for each certain attack (a, b) stemming from a: $\mathbf{A}_1$: make δ_a attack b (basically, (δ_a, b) replicates the role of (a, b)); $\mathbf{A}_2$: make b counterattack a; $\mathbf{A}_3$: make every argument attacking a attack δ_a; $\mathbf{A}_4$: discard the uncertain attacks (y, a). The intuition is that the counterattack (b, a) allows b to be put in an extension provided that at least one other argument defending it against the attack from δ_a is in the extension as well, and this is the same as what happens in the original iAAF for b, that can be included in the possible extensions containing at least one argument counterattacking (a, b).

Algorithm 1 implements these intuitions as follows. Lines 3-6 transform the input IF into an iAAF IF' with a "reduced amount of uncertainty": basically, they follow intuitions $\mathcal{I}_1, \mathcal{I}_2$, as they 1) materialize the counterattacks against certain attacks from uncertain arguments, and 2) transform uncertain arguments into certain arguments. Then, the **for** loop starting at line 7 accomplishes actions $\mathbf{A}_1$, $\mathbf{A}_2$, $\mathbf{A}_3$, according to $\mathcal{I}_4$; finally, the deletions of the uncertain attacks

Algorithm 1: Translating an iAAF into an equivalent AAF under $\sigma = \texttt{ad}$

1 **Input:** An iAAF $IF = \langle A, A^?, D, D^? \rangle$
2 **Output:** An AAF F such that $IF \equiv^\sigma F$ under $\sigma = \texttt{ad}$
3 $D^* \leftarrow \{(b,a) | a \in A^? \wedge b \in A \cup A^? \wedge (a,b) \in D\}$;
4 $D' \leftarrow D \cup D^*$; $\quad D'^? \leftarrow D^? \setminus D^*$;
5 $A' \leftarrow A \cup A^?$;
6 $IF' = \langle A', \emptyset, D', D'^? \rangle$;
7 **for** $a \in A'$ *such that* $\exists (y,a) \in D'^? \wedge \exists (a,b) \in D'$ **do**
8   Let δ_a be a new argument;
9     $A' \leftarrow A' \cup \{\delta_a\}$; $\quad D' \leftarrow D' \cup \{(\delta_a, \delta_a)\}$;
10   **for** $(a,b) \in D'$ **do**
11         $D' \leftarrow D' \cup \{(\delta_a, b), (b, a)\}$;
12   **for** $(x,a) \in D' \cup D'^?$ **do**
13         $D' \leftarrow D' \cup \{(x, \delta_a)\}$;
14 **return** the AAF $F = \langle A', D' \rangle$;

(suggested by $\mathcal{I}_3$ and $\mathcal{I}_4$) are implicitly performed by not including the uncertain attacks in the output AAF (line 14).

Performing these steps over the iAAFs on the left-hand side of Figure 2 yields the AAFs on its right-hand side. As for IF': 1) the attack (z,a) (stemming from the uncertain argument z) yields the counterattack (a,z); 2) all the uncertain arguments are made certain; 3) the uncertain attacks (y,a) and (b,a) are removed and trigger no other action, as their target argument is no source of certain attacks. As for IF, the attack (z,a) again yields the counterattack (a,z), while the chain of attacks $(y,a), (a,b)$ is expanded into the portion of AAF consisting of the self-contradicting argument δ_a that is attacked by every argument attacking a and that attacks b (which, in turn, defends itself against a). The uncertain attack (y,a) is not included in F.

The correctness and the complexity of Algorithm 1 (which implies the size of the AAF into which an input iAAF is translated) are formally stated in what follows.

Theorem 1. *Algorithm 1 runs in linear time in the size of the input iAAF IF, and returns an AAF F such that $IF \equiv^{ad} F$.*

We now focus on the other Dungean semantics, and show a negative result. That is, it is no more always possible to transform an iAAF into an equivalent AAF, even if no limits are imposed on the size of the AAF and even if the uncertainty is limited to involve only the arguments or the attacks.

Theorem 2. *Under $\sigma \in \{\texttt{st}, \texttt{co}, \texttt{gr}, \texttt{pr}\}$, there are at least one arg-iAAF IF and one att-iAAF IF' such that no AAF F is such that $IF \equiv^\sigma F$ and no AAF F' is such that $IF' \equiv^\sigma F'$*

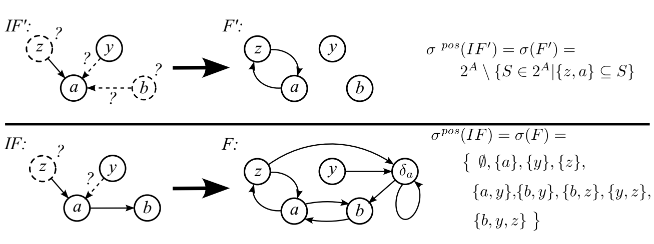

Figure 2: Examples of transformations under $\sigma = \texttt{ad}$

Proof. Under $\sigma = \texttt{co}$, the att-iAAF $IF' = \langle A, D, D^? \rangle$, with $A = \{a,b,c\}$, $D = \{(b,c)\}$, $D^? = \{(a,b)\}$, and the arg-iAAF $IF = \langle A, A^?, D \rangle$, with $A = \{b,c\}$, $A^? = \{a\}$, $D = \{(a,b),(b,c)\}$, admit no equivalent AAF, since both IF' and IF have two possible extensions whose intersection is not a possible extension, while intersecting any two complete extensions of an AAF yields a complete extension.

As for $\sigma \in \{\texttt{st}, \texttt{pr}, \texttt{gr}\}$, counterexamples proving the statements are the att-iAAF $IF' = \langle A, D, D^? \rangle$, with $A = \{a,b\}$, $D = \emptyset$, $D^? = \{(a,b)\}$, and the arg-iAAF $IF = \langle A, A^?, D \rangle$, with $A = \{a\}$, $A^? = \{b\}$, $D = \emptyset$. The point is that in every AAF F: 1) under $\sigma \in \{\texttt{st}, \texttt{pr}\}$, $\nexists S', S'' \in \sigma(F)$ such that $S' \subset S''$; 2) under $\sigma = \texttt{gr}$, $|\sigma(F)| = 1$. $\square$

Theorem 2 does not rule out the existence of a translation from iAAFs to projection-equivalent AAFs under Dungean semantics different from ad. We show that this holds for $\sigma \in \texttt{st}, \texttt{pr}$, starting with the stable semantics.

We first provide the intuitions that will guide us towards the definition of Algorithm 2, that transforms iAAFs into projection-equivalent AAFs under $\sigma = \texttt{st}$:

$\mathcal{I}_1$: under the stable semantics, every extension S has this property: for each argument a, either a is in S, or an argument x_a attacking it is in S. Hence, in an iAAF, an uncertain argument a can be transformed into a certain argument in mutual attack with a new argument x_a: the possible extensions not containing a will become possible extensions containing x_a, while those containing a are preserved;

$\mathcal{I}_2$: once all the uncertain arguments have been made certain as described above, the only uncertain components still occurring in the iAAF (if any) are uncertain attacks between certain arguments. Thus, a single uncertain attack (a,b) can be simulated by replacing it with a new certain portion of iAAF, consisting of two new arguments δ'_{ab}, δ''_{ab}, and the attacks (a, δ'_{ab}), $(\delta'_{ab}, \delta''_{ab})$, (δ''_{ab}, b), (b, δ''_{ab}) (see Figure 3). In fact, there is a one-to-one correspondence between the sets of possible extensions before and after the replacement of (a,b). Specifically, every set S that was a possible extension before the replacement is still a possible extension, if both $a \in S$ and $b \in S$; otherwise, S gives rise to the possible extension S' resulting from augmenting S with δ'_{ab}, if $a \notin S$, or with δ''_{ab}, if $a \in S \wedge b \notin S$.

Following these intuitions, Algorithm 2 is divided in two phases: lines 3-7 transform every uncertain argument a into a certain argument with a mutual attack with the new argument x_a. Then, lines 9-12 create, for each uncertain attack (a,b), the new arguments $\delta'_{ab}, \delta''_{ab}$ and the above described attacks involving them and a, b. Finally, Algorithm 2 inserts the new sets of arguments and attacks A^* and D^* obtained this way in the returned AAF, where the uncertain attacks of the input iAAF are not included. Figure 3 shows an example of transformation made by Algorithm 2, along with the set of possible extensions of the input iAAF and the set of extensions of the returned projection-equivalent AAF.

We now address the preferred semantics. The main idea underlying our strategy for transforming an iAAF into a projection-equivalent AAF is based on forcing the reasoning over the returned AAF to consider the completions of the original iAAF. More specifically, the following intuitions will

Algorithm 2: Translating an iAAF into a projection-equivalent AAF under $\sigma = \mathtt{st}$

1 **Input:** An iAAF $IF = \langle A, A^?, D, D^? \rangle$
2 **Output:** An AAF F such that $IF \equiv^{\sigma}_{\Pi} F$ under $\sigma = \mathtt{st}$
3 $\hat{A} \leftarrow \{x_a | a \in A^?\}$;
4 $A' \leftarrow A \cup A^? \cup \hat{A}$;
5 $\hat{D} \leftarrow \{(x_a, a), (a, x_a) | a \in A^?\}$;
6 $D' \leftarrow D \cup \hat{D}$;
7 $IF' = \langle A', \emptyset, D', D^? \rangle$;
8 $A^* = \emptyset$; $D^* = \emptyset$;
9 **for** $(a, b) \in D^?$ **do**
10 Let δ'_{ab} and δ''_{ab} be new arguments;
11 $A^* \leftarrow A^* \cup \{\delta'_{ab}, \delta''_{ab}\}$;
12 $D^* \leftarrow D^* \cup \{(a, \delta'_{ab}), (\delta'_{ab}, \delta''_{ab}), (\delta''_{ab}, b), (b, \delta''_{ab})\}$;
13 **return** the AAF $F = \langle A' \cup A^*, D' \cup D^* \rangle$;

guide the transformation:
$\mathcal{I}_1$: if we focus on the uncertain arguments, the various completions of the original iAAF come from choosing, for each of these arguments, if it is included or not. For each uncertain argument a, this choice can be simulated by making it certain and then augmenting the iAAF with a pair of arguments x_a, $\overline{x}_a$ in mutual attack, along with the attack $(\overline{x}_a, a)$. This way, a possible extension S coming from a completion C in the original iAAF corresponds to a possible extension S' of the new iAAF containing x_a if $a \in C$, and $\overline{x}_a$ if $a \notin C$.
$\mathcal{I}_2$: once all the uncertain arguments have been made certain, the various completions come from choosing, for each of uncertain attack (a, b), if it is included or not. Again, this choice can be simulated by means of a pair of arguments δ_{ab}, $\overline{\delta}_{ab}$ in mutual attack, that are connected to the chain (a, δ'_{ab}), $(\delta'_{ab}, \delta''_{ab})$, (δ''_{ab}, b) by means of the attack $(\overline{\delta}_{ab}, \delta''_{ab})$. Such a chain introduces two new arguments δ'_{ab}, δ''_{ab} and plays a role similar to the chain introduced by Algorithm 2 under $\sigma = \mathtt{st}$.

Algorithm 3 implements these intuitions, as lines 3-7 make certain every uncertain argument a (and create the new portion of AAF involving the new arguments x_a, $\overline{x}_a$). Then, lines 9-12 replace each uncertain attack with the structure described in the discussion of intuition $\mathcal{I}_2$. Figure 4 shows an example of transformation made by Algorithm 3. The reader can easily check that the projection of $\sigma(F)$ over the original arguments coincides with $\sigma^{\text{pos}}(IF)$.

The correctness and the complexity of the transformation algorithms under the stable and preferred semantics are formally characterized in the following statement.

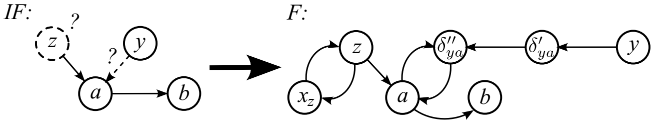

$\sigma^{pos}(IF)$:
$\{\{y, z, b\}, \{y, b\}, \{y, a\}\}$

$\sigma(F)$:
$\{\{y, z, b, \delta''_{ya}\}, \{y, b, \delta''_{ya}, x_z\}, \{y, a, x_z\}\}$

Figure 3: From an iAAF to a projection-equivalent AAF ($\sigma = \mathtt{st}$)

Algorithm 3: Translating an iAAF into a projection-equivalent AAF under $\sigma = \mathtt{pr}$

1 **Input:** An iAAF $IF = \langle A, A^?, D, D^? \rangle$
2 **Output:** An AAF F such that $IF \equiv^{\sigma}_{\Pi} F$ under $\sigma = \mathtt{pr}$
3 $\hat{A} \leftarrow \{x_a, \overline{x}_a | a \in A^?\}$;
4 $A' \leftarrow A \cup A^? \cup \hat{A}$;
5 $\hat{D} \leftarrow \{(x_a, \overline{x}_a), (\overline{x}_a, x_a), (\overline{x}_a, a) | a \in A^?\}$;
6 $D' \leftarrow D \cup \hat{D}$;
7 $IF' = \langle A', \emptyset, D', D^? \rangle$;
8 $A^* = \emptyset$; $D^* = \emptyset$;
9 **for** $(a, b) \in D^?$ **do**
10 Let δ_{ab}, $\overline{\delta}_{ab}$, δ'_{ab} and δ''_{ab} be new arguments;
11 $A^* \leftarrow A^* \cup \{\delta_{ab}, \overline{\delta}_{ab}, \delta'_{ab}, \delta''_{ab}\}$;
12 $D^* \leftarrow D^* \cup \{(\delta_{ab}, \overline{\delta}_{ab}), (\overline{\delta}_{ab}, \delta_{ab}), (\overline{\delta}_{ab}, \delta''_{ab}), (a, \delta'_{ab}), (\delta'_{ab}, \delta''_{ab}), (\delta''_{ab}, b)\}$;
13 **return** the AAF $F = \langle A' \cup A^*, D' \cup D^* \rangle$;

Theorem 3. *Algorithms 2 and 3 run in linear time in the size of the input iAAF IF, and return an AAF F such that $IF \equiv^{\sigma}_{\Pi} F$, with $\sigma = st$ and $\sigma = pr$, respectively.*

The theorem below states that what happens under $\sigma \in \{\mathtt{st}, \mathtt{pr}\}$ cannot be extended to $\sigma \in \{\mathtt{co}, \mathtt{gr}\}$, even if no polynomial bound is imposed on the size of the AAF.

Theorem 4. *Under $\sigma \in \{co, gr\}$, there is at least one arg-iAAF IF and one att-iAAF IF' such that there is no AAF F such that $IF \equiv^{\sigma}_{\Pi} F$ and no AAF F' such that $IF' \equiv^{\sigma}_{\Pi} F'$.*

5 Properties and Usability of the Transformations

By looking into the algorithms, it is easy to see that the transformations satisfy the following properties:

– *linearity of the size*: the sizes of the AAFs returned by Algorithms 1, 2, 3 are linear in the size of the input iAAF (as implied by that fact that the algorithms run in linear time);
– *preservation of arguments and of certain attacks*: all the arguments and the certain attacks in the original iAAF still occur in the returned AAF;
– *argument-invariance (only under $\sigma = ad$)*: when the uncertainty involves only arguments and/or attacks toward arguments attacking no further argument, the transformation introduces no new arguments (besides preserving all the original arguments and the certain attacks, as said above).

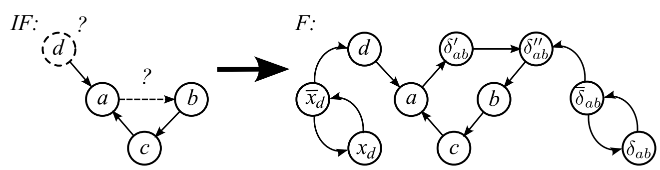

$\sigma^{pos}(IF)$:
$\{\emptyset, \{d, b\}, \{a, b\}\}$

$\sigma(F)$: $\{\{\overline{x}_d, \delta_{ab}\}, \{x_d, \delta_{ab}, \delta'_{ab}, d, b\}, \{\overline{x}_d, \overline{\delta}_{ab}, a, b\}, \{x_d, \overline{\delta}_{ab}, \delta'_{ab}, d, b\}\}$

Figure 4: From an iAAF to a projection-equivalent AAF ($\sigma = \mathtt{pr}$)

Overall, these properties say that the transformations generate AAFs where a portion of the original argumentation graph is preserved and whose size does not explode. In the case where the argument invariance is guaranteed, no new dummy arguments are needed, so the resulting AAF may be even easily readable by expert analysts. In any case, in what follows we study how the obtained AAFs, even if they did not prove easily interpretable by humans, are useful to support the reasoning over the original iAAF.

Proposition 1. *Let $IF = \langle A, A^?, D, D^? \rangle$ be an iAAF, S a set of its arguments, and F^{ad}, F^{st}, F^{pr} the AAFs returned by running Algorithms 1, 2, 3 over IF, respectively. Then:*

- *under $\sigma = \texttt{ad}$, $S \in \sigma^{pos}(IF)$ iff $S \in \sigma(F^{ad})$;*
- *under $\sigma = \texttt{st}$, $S \in \sigma^{pos}(IF)$ iff $S' \in \sigma(F^{st})$, where: $S' = S \cup \{x_a | a \in A^? \wedge a \notin S\} \cup \{\delta'_{ab} | (a,b) \in D^? \wedge a \notin S\} \cup \{\delta''_{ab} | (a,b) \in D^? \wedge a \in S \wedge b \notin S\}$;*
- *under $\sigma = \texttt{pr}$, $S \in \sigma^{pos}(IF)$ iff $\exists S' \in \sigma(F^{pr})$ such that:*
  1) *$S = \Pi_{A \cup A^?}(S')$, and*
  2) *for each $a \in A^?$, it holds that: if $a \in S$ then $x_a \in S'$, otherwise, exactly one between x_a and $\overline{x}_a$ is in S';*
  3) *for each $(a,b) \in D^?$ it holds that: $\Pi_{\{\delta_{ab}, \overline{\delta}_{ab}, \delta'_{ab}, \delta''_{ab}\}}(S') \in \big\{\{\overline{\delta}_{ab}\}, \{\overline{\delta}_{ab}, \delta'_{ab}\}, \{\delta_{ab}\}, \{\delta_{ab}, \delta'_{ab}\}, \{\delta_{ab}, \delta''_{ab}\}\big\}$.*
- *under $\sigma \in \{\texttt{ad}, \texttt{st}, \texttt{pr}\}$, $\textsc{pca}^\sigma(IF) = \Pi_{A \cup A^?}\big(\textsc{ca}^\sigma(F^\sigma)\big)$;*
- *under $\sigma \in \{\texttt{co}, \texttt{pr}\}$, $\textsc{pca}^\sigma = \Pi_{A \cup A^?}\big(\textsc{ca}^{ad}(F^{ad})\big)$;*
- *under $\sigma \in \{\texttt{st}, \texttt{pr}\}$, $\textsc{nsa}^\sigma(IF) = \Pi_A\big(\textsc{sa}^\sigma(F^\sigma)\big)$.*

Basically, the above proposition states that there is a linear-time computable bijection f between $\sigma^{\text{pos}}(IF)$ and $\sigma(F^\sigma)$ under $\sigma \in \{\texttt{ad}, \texttt{st}\}$. The correspondence in the left-to-right direction implies that the problem of verifying a specific set S as a possible extension of IF is equivalent to verifying $f(S)$ as an extension of F^σ, thus providing an efficient scheme for solving the verification over iAAFs by resorting to standard verification solvers over AAFs. On the other hand, the right-to-left correspondence allows us to solve the enumeration and the count of the possible extensions over iAAFs by invoking solvers of these problems over AAFs. Under the preferred semantics, instead, the above proposition implies that each possible extension of IF may correspond to an exponential number of extensions of F^{pr}. For instance, in the case in Figure 4, the possible extension $S = \{d, b\}$ corresponds to the two extensions $\{x_d, \delta_{ab}, \delta'_{ab}, d, b\}$ and $\{x_d, \overline{\delta}_{ab}, \delta'_{ab}, d, b\}$ (in the sense that both the projections of these extensions over $A \cup A^?$ are S). The reason is that S is an extension in two completions of IF (those containing d, and differing in the presence/absence of (a, b)), so it corresponds to two extensions of F^{pr}. In fact, F^{pr} encodes the completions of IF via the pairs of arguments x_a, $\overline{x}_a$ and δ_{ab}, $\overline{\delta}_{ab}$, so the extensions of F^{pr} corresponding to S both contain x_d and differ in the presence of δ_{ab} or $\overline{\delta}_{ab}$. Hence, checking if a set is a possible extension of IF may require multiple invocations of extension-checkers over F^{pr} (that is: checking if $S \in \sigma^{\text{pos}}(IF)$ requires to check if at least one S' augmenting S as stated in Proposition 1 is in $\sigma(F^{pr})$). Vice versa, the fact that multiple extensions of F^{pr} may correspond to the same possible extension of IF makes any machinery counting the extensions of F^{pr} not usable for counting the possible extensions of IF. However, enumeration algorithms for preferred extensions can be still used on F^{pr} for obtaining (after projections) the possible extensions of IF: in this regard, it should be studied (experimentally and theoretically) when enumerating on F^{pr} and then projecting is more efficient than enumerating the completions of IF, and then enumerating the extensions of each completion and removing the duplicates (as a set can be extension in different completions).

As for the acceptance, Proposition 1 ensures that running traditional solvers for the credulous and skeptical acceptance over the AAFs returned by algorithms 1, 2, and 3 allows us to solve the possible credulous acceptance under $\sigma \in \{\texttt{ad}, \texttt{st}, \texttt{pr}\}$ and the necessary skeptical acceptance under $\sigma \in \{\texttt{st}, \texttt{pr}\}$, without the need to resort to an ad hoc machinery for these acceptances over iAAFs (in particular, $\textsc{pca}^{\text{co}}$ and $\textsc{pca}^{\text{pr}}$ reduce to the case $\sigma = \texttt{ad}$). An analogous result does not hold for PSA and NCA: checking if an argument is PSA or NCA requires reasoning not only on the possible extensions, but also on the completions where these extensions come from, since these acceptances encode the membership to one extension of each completion or to every extension of one completion, respectively. In fact, as for $\sigma \in \{\texttt{ad}, \texttt{st}\}$, the level of completions is not encoded at all in the AAFs F^{ad} and F^{st}: indeed, the bijection between $\sigma^{\text{pos}}(IF)$ and $\sigma(F^\sigma)$ does not allow to reason on the multiple completions of IF where a possible extension S comes from by only looking into S as an extension of F^σ. As for $\sigma = \texttt{pr}$, as explained above, the AAF F^{pr} instead encodes the completions of IF. However, the CA and SA tests for an argument a do not allow to impose complex conditions on the arguments occurring with a in the extensions (in particular, on the arguments encoding the completions), so they cannot be used over F^{pr} to solve PSA and NCA.

Remark on $\sigma = \texttt{pr}$. We have discussed how the results of Proposition 1 limit the re-usability of standard machinery for the verification problem over AAFs for the case of iAAFs. These limitations would be overcome if we could resort to a new transformation from iAAFs to projection-equivalent AAFs which admits an efficiently computable function associating the possible extensions of the input iAAF to the extensions of the returned AAF. Proposition 2 states that this is not possible: due to the strictly higher complexity of verification problems for iAAFs compared to AAFs, no such transformation can exist.

Proposition 2. *Under the standard assumption $\text{P}^{\text{NP}} \neq \text{NP}^{\text{NP}}$ and under $\sigma = \texttt{pr}$, there is no polynomial-time transformation from any iAAF IF to a projection-equivalent AAF F such that the function f defined as:*

$$\forall S \in \sigma^{pos}(IF)\ f(S) = \{S' \in \sigma(F) | S = \Pi_{A \cup A^?} S'\}$$

is computable in polynomial time.

6 Related Work

Most of the research works dealing with iAAFs use the notion of extension addressed in this work (some other works adopt a different semantics, where the notions of conflict-freeness and defenses are adapted to account for the presence of uncertainty [Coste-Marquis *et al.*, 2007; Cayrol *et al.*, 2007;

Mailly, 2023; Mailly, 2024]). In particular, the possible extensions studied in this work were introduced in [Fazzinga *et al.*, 2020; Fazzinga *et al.*, 2026] (where they are called i*-extensions) to fix some counterintuitive behavior exhibited by the first proposal of the notion extension to iAAFs [Baumeister *et al.*, 2018], and are at the basis of several generalizations of iAAFs (such as [Fazzinga *et al.*, 2021a; Fazzinga *et al.*, 2021b] and [Fazzinga *et al.*, 2023] where the possibility of specifying correlations between uncertain arguments/attacks and of specifying supports between arguments have been studied, respectively). A first attempt to reduce the incompleteness in iAAFs was accomplished in [Alfano *et al.*, 2023], where it was shown that, under $\sigma \in \{\texttt{st}, \texttt{gr}, \texttt{co}, \texttt{pr}\}$, an iAAF $IF = \langle A, A^?, D, D^? \rangle$ can be transformed into an arg-iAAF or att-iAAF IF' such that $\sigma^{\text{pos}}(IF) = \Pi_{A \cup A^?} \sigma^{\text{pos}}(IF')$. Our results under $\sigma \in \{\texttt{st}, \texttt{pr}\}$ are stronger, as our transformations completely remove the incompleteness. Moreover, in some sense, our analysis strengthen the results in [Alfano *et al.*, 2023] under $\sigma \in \{\texttt{gr}, \texttt{co}\}$, as we prove that it is not possible to completely remove the incompleteness. Moreover, under $\sigma = \texttt{ad}$, we also provide the strong and elegant result regarding the equivalence between iAAFs and AAFs.

The results in this paper complete and put order in the picture of iAAFs provided by characterization of the computational complexity of the (possible) verification and acceptance problems in [Baumeister *et al.*, 2021; Fazzinga *et al.*, 2020; Fazzinga *et al.*, 2026]. In this regard, in [Fazzinga *et al.*, 2020], the (possible) verification problem over iAAFs has been shown to be in P under $\sigma \in \{\texttt{ad}, \texttt{st}, \texttt{gr}, \texttt{co}\}$ (the same as over AAFs) and Σ_p^2-complete under $\sigma = \texttt{pr}$ (that is one level higher than over AAFs, where it is coNP-complete). Basically, our results show that there is no one-to-one correspondence between the equality of complexity of the verification problem over iAAFs and AAFs and the existence of (projection-)equivalent transformations from iAAFs and AAFs. In fact, under the semantics where the complexities coincide, it can happen that: 1) an equivalent transformation exists (case $\sigma = \texttt{ad}$), 2) an equivalent transformation does not exist, but a projection-equivalent transformation exists (case $\sigma = \texttt{st}$), 3) no projection-equivalent transformation exists (case $\sigma \in \{\texttt{co}, \texttt{gr}\}$). Then, when we turn our attention to $\sigma = \texttt{pr}$, under which the verifications over iAAFs and AAFs are not of the same complexity, again it happens that the existence of a projection-equivalent AAF for any given iAAF is guaranteed. Observe that under $\sigma \in \{\texttt{ad}, \texttt{st}\}$, Algorithms 1 and 2 can be used as reductions from the verification problem over an iAAF IF to the verification problem over an AAF F, since in both cases they build AAFs with a one-to-one correspondence between $\sigma^{\text{pos}}(IF)$ and $\sigma(F)$. As a matter of fact, the existence of such a transformation/reduction is compatible with the equality of the complexities of the verifications over iAAFs and over AAFs under $\sigma \in \{\texttt{ad}, \texttt{st}\}$. On the other hand, under $\sigma = \texttt{pr}$, since verification over iAAFs is strictly harder than over AAFs, Proposition 2 guarantees that no such transformation/reduction exists.

This work is also related to the study of equivalence between AAFs [Baumann, 2012; Baumann *et al.*, 2017; Dung, 1995; Oikarinen and Woltran, 2010] and realizability [Baumann and Heine, 2023; Dunne *et al.*, 2015]. As for the equivalence, our notion of equivalence extends to iAAFs the so called "standard equivalence" between AAFs. Projection-equivalence has no corresponding notion of equivalence between AAFs in the literature cited above. As for the realizability, in [Dunne *et al.*, 2015] the signatures of Dungean semantics have been addressed (the signature of a semantics σ identifies the sets of sets of arguments Σ_σ that can be "realized" as an AAF under σ, that is $\Sigma_\sigma = \{\sigma(F) | F$ is an AAF $\}$). It is easy to see that $\Sigma_{\texttt{ad}}$ is closed under union. This implies that, since an iAAF can be seen as a set of AAFs, any iAAF always admits an equivalent AAF. However, this says nothing about the composition of this AAF (which might be of exponential size) and about a strategy for constructing it. Our result (Theorem 1) is much stronger, as we show a linear-time algorithm yielding an equivalent AAF. Under $\sigma \in \{\texttt{st}, \texttt{pr}\}$, it can be seen that Σ_σ is not closed under union. As a matter of fact, our proof of Theorem 2 for these semantics use counterexamples consisting of iAAFs that do not conform to the signatures. Thus, our results on the existence of linear-time transformations of iAAFs into projection equivalent AAFs go far beyond the negative results derivable by this analysis of the signatures.

7 Conclusions and Future Work

We have addressed a gap in the literature on (incomplete) abstract argumentation by identifying the Dungean semantics under which an iAAF can be transformed into an equivalent or projection-equivalent AAF. Beyond clarifying the actual expressive contribution of allowing uncertainty over the components of an AAF, our results show that, in the cases where transformations exist, reasoning tasks for iAAFs can be reduced to reasoning over standard AAFs, thereby enabling the use of well-established argumentation machinery. At the same time, this does not undermine the role of iAAFs: even when such transformations are available, iAAFs remain a natural and human-readable formalism for the qualitative specification of uncertainty and for the compact representation of the alternative scenarios induced by incompleteness.

Future work includes several directions: extending the study to the necessary perspective; identifying heuristics to simplify transformations in specific cases; experimenting with reasoning over iAAFs by applying our transformations to AAFs and benchmarking against approaches natively designed for iAAFs [Baumeister *et al.*, 2021; Fazzinga *et al.*, 2020]. Finally, the implications for explainable AI are worth investigating, particularly whether explanation methods developed for standard AAFs remain interpretable when applied to AAFs generated from iAAFs.

References

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

[Baumann and Heine, 2023] Ringo Baumann and Anne-Marie Heine. On Conflict-free Labellings – Realizability, Construction and Patterns of Redundancy. In *Proc.*

Int. Conf. on Principles of Knowledge Representation and Reasoning (KR), 2023.

[Baumann *et al.*, 2017] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, and S. Woltran. A general notion of equivalence for abstract argumentation. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2017.

[Baumann, 2012] Ringo Baumann. Normal and strong expansion equivalence for argumentation frameworks. *Artif. Intell.*, 193, 2012.

[Baumeister *et al.*, 2018] Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack. Verification in incomplete argumentation frameworks. *Artificial Intelligence*, 264, 2018.

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

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

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

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

[Dunne *et al.*, 2015] Paul Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence*, 228:153–178, 2015.

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

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

[Fazzinga *et al.*, 2020] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2020.

[Fazzinga *et al.*, 2021a] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over argument-incomplete aafs in the presence of correlations. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2021.

[Fazzinga *et al.*, 2021b] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over attack-incomplete aafs in the presence of correlations. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR)*, 2021.

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

[Fazzinga *et al.*, 2023] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Incomplete bipolar argumentation frameworks. In *Prof. European Conf. on Artificial Intelligence (ECAI)*, 2023.

[Fazzinga *et al.*, 2026] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Revisiting the notions of extension and acceptance over incomplete abstract argumentation frameworks. *Artif. Intell.*, 354, 2026.

[Governatori *et al.*, 2004] Guido Governatori, Michael J. Maher, Grigoris Antoniou, and David Billington. Argumentation semantics for defeasible logic. *J. Log. Comput.*, 14(5):675–702, 2004.

[Herzig and Yuste Ginel, 2021] Andreas Herzig and Antonio Yuste Ginel. Multi-agent abstract argumentation frameworks with incomplete knowledge of attacks. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1922–1928, 2021.

[Li *et al.*, 2011] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proc. Work. Theorie and Applications of Formal Argumentation (TAFA)*, 2011.

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

[Mailly, 2024] Jean-Guy Mailly. Grounded semantics and principle-based analysis for incomplete argumentation frameworks. *Int. J. Approx. Reason.*, 175, 2024.

[Oikarinen and Woltran, 2010] Emilia Oikarinen and S. Woltran. Characterizing strong equivalence for argumentation frameworks. *Artif. Intell.*, 175, 2010.