Non-flat ABA Is an Instance of Bipolar Argumentation

Markus Ulbricht[1], Nico Potyka[2], Anna Rapberger[3], Francesca Toni[3]

[1]ScaDS.AI, Department of Computer Science, Leipzig University
[2]School of Computer Science and Informatics, Cardiff University
[3]Department of Computing, Imperial College London
mulbricht@informatik.uni-leipzig.de, PotykaN@cardiff.ac.uk, {f.toni,a.rapberger}@imperial.ac.uk

Abstract

Assumption-based Argumentation (ABA) is a well-known structured argumentation formalism, whereby arguments and attacks between them are drawn from rules, defeasible assumptions and their contraries. A common restriction imposed on ABA frameworks (ABAFs) is flatness, that is, each of the defeasible assumptions can only be assumed, but not derived. While it is known that flat ABAFs can be translated into abstract argumentation frameworks (AFs) as proposed by Dung, no translation exists from general, possibly non-flat ABAFs into any kind of abstract argumentation formalism. In this paper, we close this gap and show that bipolar AFs (BAFs) can instantiate general ABAFs. To this end we develop suitable, novel BAF semantics which borrow from the notion of deductive support. We investigate basic properties of our BAFs, including computational complexity, and prove the desired relation to ABAFs under several semantics.

1 Introduction

Computational models of argumentation (Baroni et al. 2018) play a central role in non-monotonic reasoning and have a wide range of applications in various fields such as law and healthcare (Atkinson et al. 2017). One key aspect of these models is the use of structured argumentation formalisms (Besnard et al. 2014), which outline formal argumentative workflows from building blocks. Prominent approaches include assumption-based argumentation (ABA) (Bondarenko et al. 1997), ASPIC$^+$ (Modgil and Prakken 2013), DeLP (García and Simari 2004), and deductive argumentation (Besnard and Hunter 2008). The reasoning process within these formalisms typically involves creating argument structures and identifying conflicts among them in a systematic manner from rule-based knowledge bases. The resulting arguments and conflicts are known as argumentation frameworks (AFs) (Dung 1995). These frameworks are then evaluated using semantics to resolve conflicts, determine the acceptability of arguments and draw conclusions based on the original knowledge bases.

In this paper, we focus on ABA, as a versatile structured argumentation formalism that, despite its simplicity, is able to handle reasoning with certain types of preferences, strict and defeasible rules, and different types of attacks without the need for additional tools (Toni 2014). Additionally, ABA can be naturally expanded to include more advanced reasoning with preferences (Cyras and Toni 2016) and probabilities (Dung and Thang 2010), can support applications (e.g. in healthcare (Craven et al. 2012; Cyras et al. 2021a), law (Dung, Thang, and Hung 2010) and robotics (Fan et al. 2016)), and can be suitably deployed in multi-agent settings to support dialogues (Fan and Toni 2014).

An ABA framework (ABAF) amounts to a set of *rules* from some deductive system, candidate *assumptions* amongst the sentences in its language, and a *contrary* for each assumption: at an abstract level, arguments are deductions supported by rules and assumptions, and attacks are directed at assumptions in the support of arguments, by means of arguments for their contraries. Thus, assumptions constitute the defeasible part of an ABAF. There are no restrictions in ABA, in general, as to where the assumptions may appear in the rules (Čyras et al. 2018). A common restriction adopted in the study and deployment of ABA, however, is that ABAFs are *flat*, i.e., each set of assumptions is *closed* (Bondarenko et al. 1997). Intuitively speaking, flatness means that assumptions cannot be inferred, only assumed to be true or not. Flat ABA is well-studied and, because of its relative simplicity (Cyras, Heinrich, and Toni 2021), it is equipped with a variety of computational mechanisms (Toni 2013; Bao, Cyras, and Toni 2017; Lehtonen, Wallner, and Järvisalo 2022). However, general, *non-flat* ABAFs have not yet been studied as comprehensively (except for aspects of computational complexity, e.g. as in (Cyras, Heinrich, and Toni 2021)), despite their potential for a broader range of applications than restricted flat ABA. The following example gives a simple illustration.

Example 1.1. *Consider the following discussion about climate change. It is an abstraction of an idealised but realistic debate: is climate change actually happening? We cannot prove it for sure, so it makes sense to see it as an assumption (cc), but we can try and establish it by looking at its consequences: if it is actually happening we may expect an increased amount of rain (assumption mr), but then may need to deal with arguments against the validity of this assumption: one may argue that there has always been more rain at times, and so it is standard (assumption sr for "standard rain") and thus object against mr (using rule $not\_mr \leftarrow sr$), which in turn can be defeated by looking at* *statistics (s). This yields an ABAF D consisting of atoms $\mathcal{L} = \{cc, mr, sr, s, not\_cc, not\_mr, not\_sr\}$, assumptions $\mathcal{A} = \{cc, mr, sr\}$, and rules ($\mathcal{R}$):*

$$mr \leftarrow cc, \quad not\_mr \leftarrow sr, \quad not\_sr \leftarrow s, \quad s \leftarrow,$$

Moreover, for each assumption X, the contrary *is not_X. This is a non-flat ABAF, as the assumption mr is derivable from the assumption cc. Allowing for assumptions to be derived from rules can thus accommodate a form of hypothetico-deductive reasoning in debates of this form.*

A main booster for the development of flat ABAFs was the close correspondence to abstract AFs (Čyras et al. 2018). This contributed to the theoretical understanding of flat ABA, but plays also an important role in further aspects like explainability (Cyras et al. 2021b), dynamic environments (Rapberger and Ulbricht 2023), and solving reasoning tasks (Lehtonen et al. 2023).

In this paper, we extend this line of research and establish a connection between non-flat ABA and an abstract argumentation formalism. To this end, we require two ingredients. The first crucial observation is that ABAFs can be translated into *bipolar* AFs (BAFs) (Karacapilidis and Papadias 2001; Cayrol and Lagasquie-Schiex 2005; Amgoud et al. 2008) under a novel semantics. As opposed to Dung-style AFs, BAFs do not only consider an *attack* relation between arguments, representing conflicts, but also a *support* relation, that can represent justifications. Various semantics for BAFs have been proposed in the literature (see (Cayrol and Lagasquie-Schiex 2013; Cohen et al. 2014) for overviews). Our BAF semantics, which capture non-flat ABAFs, borrow ideas from previous approaches, but are novel in their technical details. The second observation is that the aforementioned approach does not work for all common ABA semantics. We tackle this issue by slightly extending our BAFs, similarly in spirit to so-called claim-augmented AFs (CAFs) (Dvorák, Rapberger, and Woltran 2020) which assign to each argument a corresponding claim. In our work, we will extend BAFs with *premises* storing under which conditions an argument can be inferred.

The main contributions of this paper are as follows.

- We define BAF semantics: novel, albeit similar in spirit to an existing semantics interpreting support as deductive (Boella et al. 2010). We also study basic properties.
- We show that for complete-based semantics, non-flat ABAFs admit a translation to BAFs w.r.t. our semantics.
- We propose so-called premise-augmented BAFs and show that they capture all common ABA semantics.
- We analyse the computational complexity of our BAFs.

2 Background

Abstract Argumentation. An abstract argumentation framework (AF) (Dung 1995) is a directed graph $F = (A, \mathrm{Att})$ where A represents a set of arguments and $\mathrm{Att} \subseteq A \times A$ models *attacks* between them. For two arguments $x, y \in A$, if $(x, y) \in \mathrm{Att}$ we say that x *attacks* y, x *is an attacker of* y, as well as x *attacks* (any set) E given that $y \in E \subseteq A$. We let $E^+_F = \{x \in A \mid E \text{ attacks } x\}$.

A set $E \subseteq A$ is *conflict-free* in F iff for no $x, y \in E$, $(x, y) \in \mathrm{Att}$. E *defends* an argument x if E attacks each attacker of x. A conflict-free set E is *admissible* in F ($E \in ad(F)$) iff it defends all its elements. A *semantics* is a function $F \mapsto \sigma(F) \subseteq 2^A$. This means, given an AF $F = (A, R)$, a semantics returns a set of subsets of A. These subsets are called σ*-extensions*. In this paper we consider so-called *admissible*, *complete*, *grounded*, *preferred*, and *stable* semantics (abbr. ad, co, gr, pr, stb). For an AF and $E \in ad(F)$, we let i) $E \in co(F)$ iff E contains all arguments it defends; ii) $E \in gr(F)$ iff E is $\subseteq$-minimal in $co(F)$; iii) $E \in pr(F)$ iff E is $\subseteq$-maximal in $co(F)$; iv) $E \in stb(F)$ iff $E^+_F = A \setminus E$.

A *bipolar argumentation framework (BAF)* is a tuple $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$, where A is a finite set of arguments, $\mathrm{Att} \subseteq A \times A$ is the *attack relation* as before and $\mathrm{Sup} \subseteq A \times A$ is the *support relation* (Amgoud et al. 2008). Given a BAF $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$, we call $F = (A, \mathrm{Att})$ the underlying AF of $\mathcal{F}$. Graphically, we depict the attack relation by solid edges and the support relation by dashed edges.

Assumption-based Argumentation. We assume a deductive system $(\mathcal{L}, \mathcal{R})$, where $\mathcal{L}$ is a formal language, i.e., a set of sentences, and $\mathcal{R}$ is a set of inference rules over $\mathcal{L}$. A rule $r \in \mathcal{R}$ has the form $a_0 \leftarrow a_1, \ldots, a_n$ with $a_i \in \mathcal{L}$. We denote the head of r by $head(r) = a_0$ and the (possibly empty) body of r with $body(r) = \{a_1, \ldots, a_n\}$.

Definition 2.1. *An ABA framework is a tuple $(\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$, where $(\mathcal{L}, \mathcal{R})$ is a* deductive system*, $\mathcal{A} \subseteq \mathcal{L}$ a non-empty set of* assumptions*, and $^- : \mathcal{A} \to \mathcal{L}$ a (total)* contrary *function.*

We say that a sentence $p \in \mathcal{L}$ is *derivable* from assumptions $S \subseteq \mathcal{A}$ and rules $R \subseteq \mathcal{R}$, denoted by $S \vdash_R p$, if there is a finite rooted labeled tree T such that the root is labeled with p, the set of labels for the leaves of T is equal to S or $S \cup \{\top\}$, and for every inner node v of T there is a rule $r \in R$ such that v is labelled with $head(r)$, the number of successors of v is $|body(r)|$ and every successor of v is labelled with a distinct $a \in body(r)$ or $\top$ if $body(r) = \emptyset$

By $Th_D(S) = \{p \in \mathcal{L} \mid \exists S' \subseteq S : S' \vdash_R p\}$ we denote the set of all conclusions derivable from an assumption-set S in an ABA framework (ABAF) D. Observe that $S \subseteq Th_D(S)$ since, by definition, each $a \in \mathcal{A}$ is derivable from $\{a\} \vdash_\emptyset a$. For $S \subseteq \mathcal{A}$, we let $\overline{S} = \{\overline{a} \mid a \in S\}$; moreover, for a derivation $S \vdash p$ we write $asms(S \vdash p) = S$ and for a set E of derivations we let $asms(E) = \bigcup_{x \in E} asms(x)$. Also, we often write $S \vdash_R p$ simply as $S \vdash p$.

A set $S \subseteq \mathcal{A}$ *attacks* a set $T \subseteq \mathcal{A}$ if for some $a \in T$ we have that $\overline{a} \in Th_D(S)$. A set S is *conflict-free*, denoted $E \in cf(D)$, if it does not attack itself. With a little notational abuse we say S attacks a if S attacks the singleton $\{a\}$.

Given $S \subseteq \mathcal{A}$, the *closure* $cl(S)$ of S is $cl(S) = Th_D(S) \cap \mathcal{A}$. With a little notational abuse we write $cl(a)$ instead of $cl(\{a\})$ whenever S is a singleton. A set $S \subseteq \mathcal{A}$ is *closed* if $S = cl(S)$. Observe that, in order for S to be non-closed, it is necessary that $\mathcal{R}$ contains a rule $a_0 \leftarrow a_1, \ldots, a_n$ s.t. $a_0 \in \mathcal{A}$, i.e., the head of the rule is an assumption.

Now we consider defense (Bondarenko et al. 1997; Čyras et al. 2018). Observe that defense in general ABAFs is only required against *closed* sets of attackers. Formally:

Definition 2.2. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ be an ABAF, $S \subseteq \mathcal{A}$ and $a \in \mathcal{A}$. We say that S* defends *a iff for each closed set T of assumptions s.t. T attacks a, we have that S attacks T; S defends itself iff S defends each $b \in S$.*

We next recall admissible, grounded, complete, preferred, and stable ABA semantics.

Definition 2.3. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ be an ABAF and $S \subseteq \mathcal{A}$ be a set of assumptions s.t. $S \in cf(S)$. We say*

- *$S \in ad(D)$ iff S is closed and defends itself;*
- *$S \in pr(D)$ iff S is $\subseteq$-maximal in $ad(D)$;*
- *$S \in co(D)$ iff $S \in ad(D)$ and is a superset of every assumption set it defends;*
- *$S \in gr(D)$ iff $S = \bigcap_{T \in co(D)} T$;*
- *$S \in stb(D)$ iff S is closed and attacks each $x \in \mathcal{A} \setminus S$.*

In this paper we stipulate that the empty intersection is interpreted as $\emptyset$, i.e., if $co(D) = \emptyset$, then $gr(D) = \emptyset$.

Example 2.4. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ be the ABAF where $\mathcal{L} = \{a, b, c, d, \overline{a}, \overline{b}, \overline{c}, \overline{d}\}$, $\mathcal{A} = \{a, b, c, d\}$, the contrary function is given as indicated, and $\mathcal{R}$ consists of rules:*

$$\overline{b} \leftarrow a. \qquad \overline{a} \leftarrow b. \qquad \overline{d} \leftarrow b. \qquad \overline{b} \leftarrow c. \qquad d \leftarrow c.$$

Let us discuss why $S = \{b\}$ is admissible in D. First of all, b is conflict-free as it does not derive $\overline{b}$. Also, b is closed, i.e., $cl(b) = \{b\}$. Regarding defense, we have that $\{a\} \vdash \overline{b}$, but also $\{b\} \vdash \overline{a}$, so the attack is defended against. Finally, c attacks b ($\{c\} \vdash \overline{b}$), but $\{c\}$ is not closed. Indeed, $cl(c) = \{c, d\}$. Since $\{b\} \vdash \overline{d}$, this attack is also defended against.

3 Closed Extensions for Bipolar AFs

Our goal is to translate non-flat ABAFs into BAFs. In this section, we develop BAF semantics which are suitable for this endeavor (under complete-based semantics for ABAFs). To this end we interpret the support relation in the spirit of the notion of *deductive support* (Boella et al. 2010), i.e., the intuitive reading is that whenever x is accepted and x supports y, then y is also accepted. While this approach borrows from the BAF literature, to the best of our knowledge the exact definitions do not coincide with any previously proposed BAF semantics. We define extension-based semantics directly on the given BAF, without re-writing it to an AF.

We start with the notion of the closure for BAFs.

Definition 3.1. *Let $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ be a BAF. Consider the operator μ defined by*

$$\mu(E) = E \cup \{a \in A \mid \exists e \in E : (e, a) \in \mathrm{Sup}\}.$$

We call $cl(E) = \bigcup_{n \geq 1} \mu^n(E)$ the closure *of E. A set $E \subseteq A$ is called* closed *if $E = cl(E)$.*

Now we introduce the basic concepts of conflict-freeness and defense underlying our semantics. As usual, a set of arguments is said to be conflict-free whenever it does not attack itself. Our notion of defense is inspired by the way it is defined for ABA (cf. Definition 2.2).

Definition 3.2. *Let $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ be a BAF. A set $E \subseteq A$ is* conflict-free *if $E \cap E^+_{\mathcal{F}} = \emptyset$; E* defends *$a \in A$ if E attacks each closed set $S \subseteq A$ which attacks a; the* characteristic function *of $\mathcal{F}$ is $\Gamma(E) = \{a \in A \mid E \text{ defends } a\}$.*

Observe that this is a weaker condition than the defense notion of AFs since we can disregard non-closed attackers.

Example 3.3. *Let $\mathcal{F}$ be the following BAF (recall that the attack relation is depicted by solid edges and the support relation by dashed edges):*

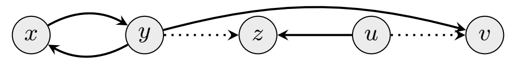

We have that $cl(y) = \{y, z\}$ and y defends z which can be seen as follows: even though u attacks z, u is not a closed set of arguments. The closure of $\{u\}$ is $cl(\{u\}) = \{u, v\}$. Since y attacks v, we find $z \in \Gamma(\{y\})$.

As we saw in this example, our defense notion can intuitively be interpreted as follows: if we seek to defend some argument a, then it suffices to counter-attack the closure of each attacker b of a (rather than b itself).

Lemma 3.4. *Let $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ be a BAF and let $E \subseteq A$ and $a \in A$. Then E defends a iff for each attacker b of a it holds that E attacks $cl(\{b\})$.*

Let us now define admissibility. We require a set of arguments to be conflict-free, closed, and self-defending.

Definition 3.5. *Let $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ be a BAF. A set $E \subseteq A$ is* admissible*, $E \in ad(\mathcal{F})$, if i) E is conflict-free, ii) E is closed, and iii) $E \subseteq \Gamma(E)$.*

Example 3.6. *Recall Example 3.3. Let us verify that $E = \{y, z\} \in ad(\mathcal{F})$. Clearly, E is closed with $E \in cf(\mathcal{F})$. The two attackers of E are x and u with $cl(x) = \{x\}$ and $cl(u) = \{u, v\}$, both of which are counter-attacked. Another admissible set is $E' = \{u, v\}$ since the attack by y is countered due to $cl(y) = \{y, z\}$ and u attacks z.*

As usual, the empty set is always admissible and hence, we can guarantee $ad(\mathcal{F}) \neq \emptyset$ for any given BAF $\mathcal{F}$.

Proposition 3.7. *Let $\mathcal{F}$ be a BAF. Then $\emptyset \in ad(\mathcal{F})$. In particular, $ad(\mathcal{F}) \neq \emptyset$.*

Given this notion of admissibility, the definition of the remaining semantics is natural: for complete extensions, we require E to include all defended arguments; preferred extensions are defined as maximal admissible sets; the grounded extension is the intersection of all complete ones.

Definition 3.8. *Let $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ be a BAF. A set $E \subseteq A$ of arguments s.t. $E \in cf(\mathcal{F})$ is*

- preferred*, $E \in pr(\mathcal{F})$, iff it is maximal admissible;*
- complete*, $E \in co(\mathcal{F})$, iff $E \in ad(\mathcal{F})$ and $E = \Gamma(E)$;*
- grounded*, $E \in gr(\mathcal{F})$, iff $E = \bigcap_{S \in co(\mathcal{F})} S$;*
- stable*, $E \in stb(\mathcal{F})$, iff it is closed and $E^+ = A \setminus E$.*

Example 3.9. *In our Example 3.3, the admissible extension $E = \{y, z\}$ is maximal and thus preferred. Moreover, $\{x, u, v\} \in pr(\mathcal{F})$. We observe however that $\{y, z\}$ is not complete since it does not contain the unattacked argument u. Hence $co(\mathcal{F}) = \{\{u, v\}\}$.*

As a final remark regarding our BAF semantics, let us mention that they do not admit a translation into Dung-style AFs. As the previous example already shows, preferred extensions are in general not complete (which is the case for

AFs). Another interesting observation is that we do not necessarily have a complete extension. These properties show that a translation to AFs is impossible.

Example 3.10. *Let $\mathcal{F}$ be the following BAF:*

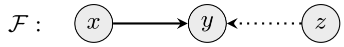

Suppose $E \in co(\mathcal{F})$. Since z is unattacked, we must have $z \in E$. As complete sets must be closed, $y \in E$ follows. However, by the same reasoning, $x \in E$ and thus, $E \notin cf(\mathcal{F})$; a contradiction. Indeed, the only admissible sets are $\emptyset$ and $\{x\}$, both of which are not complete.

Note that, for the same reason, ABAFs cannot be translated into AFs: general ABAFs violate many properties that hold for AFs. Thus we require BAFs for the translation.

4 Instantiated BAFs

Suppose we are given an ABAF $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$. Our goal is to translate D into a BAF $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$. The underlying idea is to define A and Att as it is done for flat ABAFs, i.e., each ABA argument $S \vdash p$ corresponds to an argument in $\mathcal{F}_D$ which attacks arguments in $\mathcal{F}_D$ corresponding to $T \vdash q$ whenever $p \in \overline{T}$. What we have left to discuss is the support relation. To this end we need to take care of the closure of sets of assumptions. More specifically, if some assumption a is in the closure of a set S, i.e., $S \vdash a$ is an argument in D, then we encode this in $\mathcal{F}_D$ using Sup.

To illustrate this, suppose we are given assumptions a, b, c and rules $r_1 : p \leftarrow a$, $r_2 : b \leftarrow a$, and $r_3 : \overline{b} \leftarrow c$. From r_2 it follows that $b \in cl(a)$, which can be encoded in our instantiated BAF as follows: since b is an assumption, there is some generic argument $\{b\} \vdash b$ for it, then the argument stemming from rule r_2, i.e., $\{a\} \vdash b$, supports the argument $\{b\} \vdash b$; hence any closed set accepting $\{a\} \vdash b$ must also accept $\{b\} \vdash b$. Including the usual attacks, this would give the following BAF (for $a \in \mathcal{A}$, we depict $\{a\} \vdash a$ by just a):

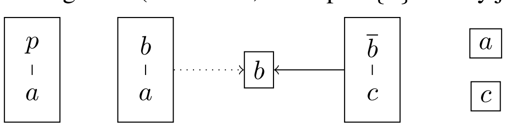

It is now indeed impossible to accept $\{a\} \vdash b$ without counter-attacking $\{c\} \vdash \overline{b}$ since the former supports $\{b\} \vdash b$. With this support relation we miss however that we cannot accept $\{a\} \vdash p$, either: since constructing this argument requires a, we would then also have to include b due to $b \in cl(a)$. Hence $\{a\} \vdash p$ should also support $\{b\} \vdash b$. More generally, an argument $S \vdash p$ shall support each $a \in cl(S)$:

$$a \in cl(S) \Rightarrow (S \vdash p, \{a\} \vdash a) \in \mathrm{Sup}. \qquad (1)$$

We therefore define the support relation of our corresponding BAF according to (1) as follows.

Definition 4.1. *For an ABAF $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$, we define the instantiated BAF $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ via*

$$A = \{(S \vdash p) \mid (S \vdash p) \textit{ is an argument in } D\}$$
$$\mathrm{Att} = \{(S \vdash p, T \vdash q) \in A^2 \mid p \in \overline{T}\}$$
$$\mathrm{Sup} = \{(S \vdash p, \{a\} \vdash a) \in A^2 \mid a \in cl(S)\}$$

Example 4.2. *Recall our ABAF D from Example 2.4. The instantiated BAF $\mathcal{F}_D$ is given as follows (again, we depict the generic argument $\{a\} \vdash a$ for each $a \in \mathcal{A}$ by a).*

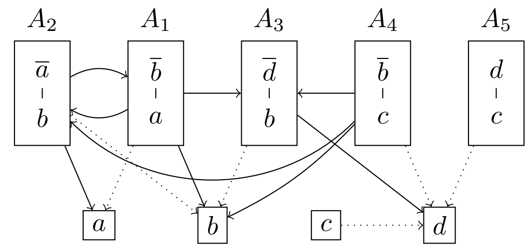

As we saw, $\{b\}$ is admissible in D. Now consider the set E of all arguments with $\{b\}$ as assumption set, i.e., $E = \{A_2, A_3, b\}$. Again, E is conflict-free and closed (in $\mathcal{F}_D$). The attack from A_1 is countered; moreover, A_4 attacks A_3; however, as A_4 supports d, the closure of A_4 (in $\mathcal{F}_D$) is $\{A_4, d\}$ with A_3 attacking d; so this attack is also countered. We infer that E is admissible in $\mathcal{F}_D$.

Though Definition 4.1 induces infinitely many arguments, they are determined by their underlying assumptions and conclusion. Hence it suffices to construct a finite BAF.

In the remainder of this section, we establish the following main result showing that the BAF $\mathcal{F}_D$ as defined above is suitable to capture non-flat ABAFs for complete-based semantics: if E is some extension (in $\mathcal{F}_D$), then a set of acceptable assumptions can be obtained by gathering all assumptions underlying the arguments in E; and if S is acceptable (in D), then all arguments constructible from the assumptions in S from a corresponding extension in the BAF $\mathcal{F}_D$.

Theorem 4.3. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. Let $\sigma \in \{co, gr, stb\}$.*

- *If $E \in \sigma(\mathcal{F}_D)$, then $asms(E) \in \sigma(D)$.*
- *If $S \in \sigma(D)$, then $\{x \in A \mid asms(x) \subseteq S\} \in \sigma(\mathcal{F}_D)$.*

If $S \in ad(D)$, then $\{x \in A \mid asms(x) \subseteq S\} \in ad(\mathcal{F}_D)$.

Example 4.4. *Recall our ABA D and $\mathcal{F}_D$ from above. As we already saw, $S = \{b\} \in ad(D)$ and indeed, $E = \{x \in A \mid asms(x) \subseteq S\} = \{A_2, A_3, b\} \in ad(\mathcal{F}_D)$ as we verified.*

As stated in the theorem, we only get one direction for the ad semantics; for the pr semantics, both directions fail. We will discuss and subsequently fix the underlying issue later.

From ABA to BAF

Translating an extension of the given ABAF into one in the BAF is the easier direction. Our first step is the following proposition which shows the desired connection between conflict-free and closed sets.

Proposition 4.5. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. Let $S \subseteq \mathcal{A}$ and let $E = \{x \in A \mid asms(x) \subseteq S\}$.*

- *If $S \in cf(D)$, then $E \in cf(\mathcal{F}_D)$.*
- *If S is closed in D, then E is closed in $\mathcal{F}_D$.*
- *If $S \in ad(D)$, then $E \in ad(\mathcal{F}_D)$.*

In order to extend this result to complete extensions, we have only left to show that all defended arguments are included in the corresponding BAF $\mathcal{F}_D$.

Proposition 4.6. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. If $S \in co(D)$, then for $E = \{x \in A \mid asms(x) \subseteq S\}$ we get $E \in co(\mathcal{F}_D)$.*

Moreover, the set of attacked assumptions is preserved. Thus, we also find stable extensions of $\mathcal{F}_D$.

Proposition 4.7. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. If $S \in stb(D)$, then for $E = \{x \in A \mid asms(x) \subseteq S\}$ we get $E \in stb(\mathcal{F}_D)$.*

Consequently, the first item in Theorem 4.3 is shown.

From BAF to ABA

Turning extensions of the instantiated BAF into extensions of the underlying ABAF is more involved and does not work for admissible sets without any further restriction. It is important to understand *why* it does not work for admissible sets since this will also demonstrate why complete-based semantics do not face this issue. The problem is related to the way we have to construct our support relation. We illustrate this in the following example.

Example 4.8. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be the ABAF where $\mathcal{L} = \{a, b, c, \overline{a}, \overline{b}, \overline{c}, p\}$, $\mathcal{A} = \{a, b, c\}$, $^-$ is as indicated, and $\mathcal{R} = \{p \leftarrow a.,\ q \leftarrow b.,\ c \leftarrow p, q.,\ \overline{c} \leftarrow c.\}$. Observe that $S = \{a, b\}$ is not admissible in D since S is not closed: indeed, we can derive p from a and q from b and thus c from S, i.e., $c \in cl(S)$. Now consider $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$:*

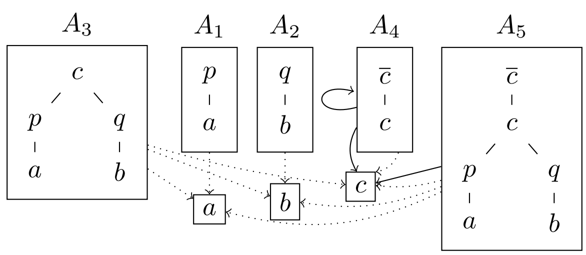

We want to emphasize that there is neither a support arrow from A_1 to c nor from A_2 to c; the fact that $c \in cl(\{a, b\})$ holds is reflected in $(A_3, c) \in \mathrm{Sup}$ and $(A_5, c) \in \mathrm{Sup}$.

Consider now $E = \{a, b, A_1, A_2\}$. As all arguments in E are unattacked and have no out-going support arrows, it is clear that $E \in ad(\mathcal{F}_D)$. Yet, the required assumptions to build these arguments are $\{a, b\}$, despite $\{a, b\} \notin ad(D)$.

The mismatch in the previous example occurred because we did not take *all* arguments we can build from a and b. Indeed, we did not include A_3 and A_5 in our extension E of $\mathcal{F}_D$. These arguments encode the support from a and b to c and, thus, we would have detected the missing c we cannot defend. This observation leads to the following notion.

Definition 4.9. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF; $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. A set $E \subseteq A$ is* assumption exhaustive *if $asms(x) \subseteq asms(E)$ implies $x \in E$.*

Example 4.10. *In the previous Example 4.8, for the set $E = \{a, b, A_1, A_2\}$ of arguments we have $asms(E) = \{a, b\}$ and hence E is not assumption exhaustive because A_3 and A_5 also satisfy $asms(A_i) \subseteq asms(E)$.*

Remark 4.11. *In the previous subsection we started with assumptions S and constructed $E = \{x \in A \mid asms(x) \subseteq S\}$. Such E is assumption exhaustive by design.*

The next proposition states that the problematic behavior we observed in Example 4.8 regarding admissible extensions does not occur for assumption exhaustive sets.

Proposition 4.12. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. Let $E \subseteq A$ be assumption exhaustive and let $S = asms(E)$.*

- *If $E \in cf(\mathcal{F}_D)$, then $S \in cf(D)$.*
- *If E is closed in $\mathcal{F}_D$, then S is closed in D.*
- *If $E \in ad(\mathcal{F}_D)$, then $S \in ad(D)$.*

Given the results we have established so far, Proposition 4.12 can only serve as an intermediate step because it relies on assumption exhaustive sets of arguments. However, within the context of an abstract BAF this notion does not make any sense; it is tailored to arguments stemming from instantiating D. Because of this, the following lemma is crucial. It states that for complete-based semantics each extension is assumption exhaustive. Thus, in this case, the mismatch we observed in Example 4.8 does not occur.

Lemma 4.13. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. If $E \in co(\mathcal{F}_D)$, then E is assumption exhaustive.*

Example 4.14. *In Example 4.8, we saw that the set $E = \{a, b, A_1, A_2\}$ of arguments is not assumption exhaustive. Indeed, since E is admissible, it is clear by definition that each argument x with $asms(x) \subseteq \{a, b\}$ is defended, including A_3 and A_5. Hence E is not complete. On the other hand, $\Gamma(E) = \{a, b, A_1, A_2, A_3, A_5\} = E'$ is assumption exhaustive as desired.*

Lemma 4.13 allows us to apply Proposition 4.12 to all complete extensions. Hence we can infer the following result without restricting to assumption exhaustive sets.

Proposition 4.15. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. If $E \in co(\mathcal{F}_D)$, then $S = asms(E) \in co(D)$.*

Corollary 4.16. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF.*

- *If $E \in gr(\mathcal{F}_D)$, then $S = asms(E) \in gr(D)$.*
- *If $S \in gr(D)$, then for $E = \{x \in A \mid asms(x) \subseteq S\}$ we have that $E \in gr(\mathcal{F}_D)$.*

Finally, stable extensions (in the BAF $\mathcal{F}_D$) are also assumption exhaustive due to Lemma 4.13 (each stable extension is complete). This yields the desired connection for stb.

Proposition 4.17. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and $\mathcal{F}_D = (A, \mathrm{Att}, \mathrm{Sup})$ the instantiated BAF. If $E \in stb(\mathcal{F}_D)$, then $S = asms(E) \in stb(D)$.*

Example 4.18. *Let us head back to our motivating Example 1.1 on climate change. We instantiate the following BAF.*

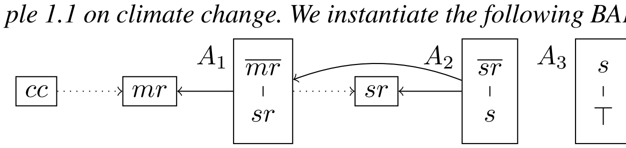

The BAF reflects the fact that cc (in favor of climate change) can only be accepted when mr (more rain) is also included in the extension. As desired, cc is acceptable; for instance, $\{cc, mr\} \in co(D)$. Correspondingly, $\{cc, mr, A_1, A_3\} \in co(\mathcal{F}_D)$ so the acceptability of cc is also found in $\mathcal{F}_D$.

5 BAFs and Admissible Semantics

Our analysis in Section 4 reveals that we cannot capture admissible (and consequently preferred) ABA semantics by means of our instantiated BAFs since there is no way to guarantee that the accepted sets of arguments are assumption exhaustive. In this section, we will propose a slightly augmented version of BAFs which additionally stores this information. This proposal is in line with recent developments in AF generalizations which capture certain features of instantiated graphs in addition to a purely abstract view (Dvorák and Woltran 2020; Rapberger 2020; Rapberger and Ulbricht 2023). In Section 6 we will see that the computational price we have to pay for this is moderate.

Definition 5.1. *Let $\mathcal{P}$ be a set (of premises). A* premise-augmented BAF *(pBAF) $\mathbb{F}$ is a tuple $\mathbb{F} = (A, \mathrm{Att}, \mathrm{Sup}, \pi)$ where $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$ is a BAF and $\pi : A \to 2^{\mathcal{P}}$ is the* premise function*; $\mathcal{F}$ is called the* underlying *BAF of $\mathbb{F}$.*

We let $\pi(E) = \bigcup_{a \in E} \pi(a)$. We sometimes abuse notation and write $\mathbb{F} = (\mathcal{F}, \pi)$ for the pBAF $\mathbb{F} = (A, \mathrm{Att}, \mathrm{Sup}, \pi)$ with underlying BAF $\mathcal{F} = (A, \mathrm{Att}, \mathrm{Sup})$. The following properties are defined due to the underlying BAF $\mathcal{F}$:

Definition 5.2. *Let $\mathbb{F} = (\mathcal{F}, \pi)$ be a pBAF. A set $E \subseteq A$ is* conflict-free *resp.* closed *whenever this is the case for E in $\mathcal{F}$; E* defends *$a \in A$ in $\mathbb{F}$ iff this is the case in $\mathcal{F}$.*

The only novel concept we require is the notion of an exhaustive set of arguments.

Definition 5.3. *Let $\mathbb{F} = (\mathcal{F}, \pi)$ be a pBAF. A set $E \subseteq A$ is* exhaustive *iff $\pi(a) \subseteq \pi(E)$ implies $a \in E$.*

Semantics for pBAFs are defined similarly as for BAFs, but with the important difference that we require all admissible-based extensions to be exhaustive.

Definition 5.4. *For a pBAF $\mathbb{F} = (\mathcal{F}, \pi)$, a set $E \in cf(\mathcal{F})$ is*

- admissible*, $E \in ad(\mathbb{F})$, iff E is exhaustive and $E \in ad(\mathcal{F})$;*
- preferred*, $E \in pr(\mathbb{F})$, iff it is $\subseteq$-maximal admissible;*
- complete*, $E \in co(\mathbb{F})$, iff $E \in ad(\mathbb{F})$ and $E = \Gamma(E)$;*
- grounded*, $E \in gr(\mathbb{F})$, iff $E = \bigcap_{S \in co(\mathbb{F})} S$;*
- stable*, $E \in stb(\mathbb{F})$, iff it is closed and $E^+_{\mathcal{F}} = A \setminus E$.*

Example 5.5. *Let us illustrate how pBAFs can help us fixing our issue illustrated in Example 4.8. Let us construct the same BAF, but assign to each argument the assumptions required to entail it as premises.*

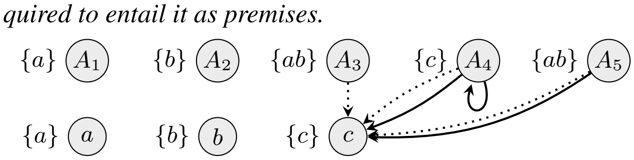

We have that, for instance, $E = \{a, A_1\} \in ad(\mathbb{F})$. Both arguments are unattacked with no out-going support arrow. Thus the only condition to verify is exhaustiveness. This property is satisfied since no further argument x satisfies $\pi(x) \subseteq \pi(E) = \{a\}$.

On the other hand, $E' = \{a, b, A_1, A_2\}$ is not admissible. Since $\pi(E') = \{a, b\}$, exhaustiveness would also require presence of A_3 and A_5 (which, in turn, would result in acceptance of c which cannot be defended).

Following the observations made in this example, we define $\mathbb{F}_D$ as follows: The underlying BAF $\mathcal{F}$ is given as before and π stores the assumptions required to entail an argument.

Definition 5.6. *For an ABAF $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$, the* instantiated pBAF *$\mathbb{F}_D = (A, \mathrm{Att}, \mathrm{Sup}, \pi) = (\mathcal{F}, \pi)$ is*

$$\mathcal{F} = \mathcal{F}_D \qquad \forall x \in A : \pi(x) = asms(x).$$

We can now capture any non-flat ABAF as follows.

Theorem 5.7. *Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ be an ABAF and $\mathbb{F}_D = (\mathcal{F}, \pi)$ the instantiated pBAF. Then*

- *if $E \in \sigma(\mathbb{F}_D)$, then $asms(E) \in \sigma(D)$;*
- *if $S \in \sigma(D)$, then $\{x \in A \mid asms(x) \subseteq S\} \in \sigma(\mathbb{F}_D)$*

for any $\sigma \in \{ad, co, pr, gr, stb\}$.

6 Computational Complexity

We consider the usual decision problems (under semantics σ) in formal argumentation. Let $\mathcal{K}$ be a knowledge base (i.e., an ABAF, BAF, or pBAF), let a be an assumption resp. argument, and let E be a set of assumptions resp. arguments.

- Credulous acceptance *Cred*$_\sigma$: Given $\mathcal{K}$ and some a, is it true that $a \in E$ for some $E \in \sigma(\mathcal{K})$?
- Skeptical acceptance *Skept*$_\sigma$: Given $\mathcal{K}$ and some a, is it true that $a \in E$ for each $E \in \sigma(\mathcal{K})$?
- Verification *Ver*$_\sigma$: Given $\mathcal{K}$ and a set E, is it true that $E \in \sigma(\mathcal{K})$?

We start with the computational complexity of BAFs with our novel semantics. The high level observation is that many tasks are close to reasoning in usual AFs. However, computing the grounded extension is much harder, inducing certain consequences (e.g. there is no shortcut for skeptical reasoning under complete semantics).

Theorem 6.1. *For BAFs, the problem*

- *Ver$_\sigma$ is tractable for $\sigma \in \{ad, co, stb\}$, coNP-complete for $\sigma = pr$, and DP-complete for $\sigma = gr$.*
- *Cred$_\sigma$ is NP-complete for $\sigma \in \{ad, co, pr, stb\}$ and DP-complete for $\sigma = gr$.*
- *Skept$_\sigma$ is trivial for $\sigma = ad$, DP-complete for $\sigma \in \{co, gr, stb\}$, and Π_2^P-complete for $\sigma = pr$.*

Surprisingly, the price we have to pay for also capturing admissible-based semantics is rather small. The computational complexity of the pBAFs we construct is almost the same; the only difference is that skeptical acceptance w.r.t. admissible semantics is not trivial anymore, but now becomes coNP-complete.

Proposition 6.2. *For pBAFs, Skept$_{ad}$ is coNP-complete.*

From this observation the following main theorem follows as a corollary of the complexity results for BAFs.

Theorem 6.3. *For pBAFs, the problem*

- *Ver$_\sigma$ is tractable for $\sigma \in \{ad, co, stb\}$, coNP-complete for $\sigma = pr$, and DP-complete for $\sigma = gr$.*
- *Cred$_\sigma$ is NP-complete for $\sigma \in \{ad, co, pr, stb\}$ and DP-complete for $\sigma = gr$.*
- *Skept$_\sigma$ is coNP-complete for $\sigma = ad$, DP-complete for $\sigma \in \{co, gr, stb\}$, and Π_2^P-complete for $\sigma = pr$.*

		ad	gr	co	pr	stb
Ver_σ	ABA	coNP-c	DP_2-c	coNP-c	Π_2^P-c	in P
	BAF	in P	DP-c	in P	coNP-c	in P
	pBAF	in P	DP-c	in P	coNP-c	in P
$Cred_\sigma$	ABA	Σ_2^P-c	DP_2-c	Σ_2^P-c	Σ_2^P-c	NP-c
	BAF	NP-c	DP-c	NP-c	NP-c	NP-c
	pBAF	NP-c	DP-c	NP-c	NP-c	NP-c
$Skept_\sigma$	ABA	Π_2^P-c	DP_2-c	DP_2-c	Π_3^P-c	DP-c
	BAF	triv.	DP-c	DP-c	Π_2^P-c	DP-c
	pBAF	coNP-c	DP-c	DP-c	Π_2^P-c	DP-c

Table 1: Complexity: Non-flat ABA vs. (p)BAFs

Table 1 summarizes our complexity results for BAFs and pBAFs and compares them to the known computational complexity of non-flat ABA (Cyras, Heinrich, and Toni 2021).[1] We observe that for each reasoning problem we consider, pBAFs are one level below non-flat ABA in the polynomial hierarchy. Moreover, BAFs and pBAFs are comparable for most reasoning problems, with skeptical reasoning for ad semantics being the only exception. This is in line with our results that pBAFs are capable of capturing admissible reasoning in ABA (cf. Theorem 5.7), whereas BAFs are not. Note that Theorem 5.7 does not contradict the differing results pertaining to computational in complexity in ABA vs. pBAFs since the instantiation procedure yields exponentially many arguments in general.

7 Discussion and Related Work

There is a rich selection of bipolar argumentation approaches in the literature (Cayrol, Cohen, and Lagasquie-Schiex 2021). The most prominent ones are deductive (Boella et al. 2010), necessary (Nouioua and Risch 2010), evidential (Oren and Norman 2008) and backing (Cohen, García, and Simari 2012) support. More recent work on classical BAFs looked at symmetry between attack and support (Potyka 2020), argument attributes (Gonzalez et al. 2021) and monotonicity (Gargouri et al. 2021).

Our notion of defense can be characterized using notions of extended attacks that occur in BAFs (Amgoud et al. 2008; Boella et al. 2010). There is a *mediated attack from a to b* if a attacks an argument c that is transitively supported by b (Boella et al. 2010). Using our notion of closure from Definition 3.2, this can be rewritten as a attacks $cl(\{b\})$. Hence, Lemma 3.4 states that E defends a iff for each attacker b of a, there is a direct or mediated attack from E to b.

Another approach is due to (Amgoud et al. 2008). Here a set of arguments S defends (Dung 1995) an argument a if S attacks every attacker of a. Further, S attacks an argument a iff there is a direct attack from S to a or S transitively supports an argument b that attacks a. This amounts to requiring that $cl(S)$ attacks a. A set of arguments S is then called *conflict free* if S does not attack any argument in S in the previous sense, that is, if $cl(S)$ does not directly attack any argument in S. They call S *admissible* iff S is conflict-free, closed under support and defends all its elements. One important difference to our notion of admissible sets is the definition of defense. While (Amgoud et al. 2008) allow defense via direct and supported attacks, we allow defense via direct or mediated attacks. For instance in $\mathcal{F} = (\{a, b, c\}, \{(a, c), (b, a)\}, \{(b, c)\})$, $\{a\}$ is admissible w.r.t. our definition because it defends itself against b via a mediated attack, is closed under support and conflict-free. However, it is not admissible w.r.t. the definition in (Amgoud et al. 2008) because here, it does not defend itself against b.

The c-admissibility in (Cayrol and Lagasquie-Schiex 2005) is close to ours, but their work uses *supported* and *indirect* defeat which is incompatible with our concepts.

(Cyras et al. 2017) also explored the relation ABAFs-BAFs, but understands the BAFs under existing semantics in terms of a restricted form of non-flat-ABAFs, rather than ABAF as BAF, under new semantics, as we do.

8 Conclusion

We translated non-flat ABAFs into BAFs. To this end we proposed novel BAF semantics to capture complete-based semantics. By means of a novel formalisms, called *premise-augmented* BAFs, we also established a correspondence between admissible-based semantics of non-flat ABA and our BAFs. We discussed basic properties of these semantics and proved the correspondence to ABA. We then investigated the computational complexity of BAFs and showed that, compared to non-flat ABA, the typical reasoning problems are one level lower in the polynomial hierarchy.

This work opens several avenues for future work. It would be interesting to extend our results to further ABA semantics like semi-stable and ideal semantics (Čyras et al. 2018), as well as a set-stable semantics (Cyras, Schulz, and Toni 2017). Further, we only discussed basic properties of our (p)BAF semantics: it would also be interesting to study a version of the fundamental lemma, existence of extensions, and the impact of restricting the graph class.

The lower computational complexity of (p)BAFs compared to non-flat ABAFs raises the question as to which extent our research can contribute to the development of efficient instantiation-based ABA solvers. As for the flat ABA instantiation by means of AFs, our constructed graphs are infinite in general. However, in the case of AFs, only finitely many arguments suffice to determine the semantics, and techniques to reduce the size of the constructed AF even further are available (Lehtonen et al. 2023).

As a final remark, most approaches for explaining reasoning in ABA construct the underlying AF and extract an argumentative explanation from it (Cyras et al. 2021b). Our research serves as the first step to enabling this machinery for non-flat ABAFs as well: the next goal would be the computation of intuitive explanations in (p)BAFs. This line of research would contribute to applications where non-flat ABAFs give natural representations, as in Example 1.1 as well as settings where agents share information (e.g. as in (Cyras, Schulz, and Toni 2017)) and may thus disagree on which information is factual and defeasible.

[1] In their paper, the standard definition of the empty intersection is used, but the results can be derived analogously.

Acknowledgements

This research was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101020934, ADIX), by J.P. Morgan and by the Royal Academy of Engineering under the Research Chairs and Senior Research Fellowships scheme, and by the Federal Ministry of Education and Research of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI. Any views or opinions expressed herein are solely those of the authors.

References

Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *Int. J. Intell. Syst.*, 23(10): 1062–1093.

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards Artificial Argumentation. *AI Magazine*, 38(3): 25–36.

Bao, Z.; Cyras, K.; and Toni, F. 2017. ABAplus: Attack Reversal in Abstract and Structured Argumentation with Preferences. In *Proc. PRIMA*, volume 10621 of *Lecture Notes in Computer Science*, 420–437. Springer.

Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.

Besnard, P.; Garcia, A.; Hunter, A.; Modgil, S.; Prakken, H.; Simari, G.; and Toni, F. 2014. Introduction to structured argumentation. *Argument Comput.*, 5(1): 1–4.

Besnard, P.; and Hunter, A. 2008. *Elements of Argumentation*. MIT Press. ISBN 978-0-262-02643-7.

Boella, G.; Gabbay, D. M.; van der Torre, L.; and Villata, S. 2010. Support in abstract argumentation. In *Proc. COMMA*, 40–51. Frontiers in Artificial Intelligence and Applications, IOS Press.

Bondarenko, A.; Dung, P. M.; Kowalski, R. A.; and Toni, F. 1997. An Abstract, Argumentation-Theoretic Approach to Default Reasoning. *Artif. Intell.*, 93: 63–101.

Cayrol, C.; Cohen, A.; and Lagasquie-Schiex, M.-C. 2021. Higher-Order Interactions (Bipolar or not) in Abstract Argumentation: A State of the Art. *Handbook of Formal Argumentation*, 2.

Cayrol, C.; and Lagasquie-Schiex, M. 2013. Bipolarity in argumentation graphs: Towards a better understanding. *Int. J. Approx. Reason.*, 54(7): 876–899.

Cayrol, C.; and Lagasquie-Schiex, M.-C. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Proc. ECSQARU'05*, volume 3571 of *LNCS*, 378–389. Springer.

Cohen, A.; García, A. J.; and Simari, G. R. 2012. Backing and Undercutting in Abstract Argumentation Frameworks. In *Proc. FoIKS*, volume 7153 of *Lecture Notes in Computer Science*, 107–123. Springer.

Cohen, A.; Gottifredi, S.; García, A. J.; and Simari, G. R. 2014. A survey of different approaches to support in argumentation systems. *Knowl. Eng. Rev.*, 29(5): 513–550.

Craven, R.; Toni, F.; Cadar, C.; Hadad, A.; and Williams, M. 2012. Efficient Argumentation for Medical Decision-Making. In *Proc. KR*. AAAI Press.

Cyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2017. Assumption-based Argumentation: Disputes, Explanations, Preferences. *FLAP*, 4(8).

Čyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2018. Assumption-Based Argumentation: Disputes, Explanations, Preferences. In *Handbook of Formal Argumentation*, chapter 7, 365–408. College Publications.

Cyras, K.; Heinrich, Q.; and Toni, F. 2021. Computational complexity of flat and generic Assumption-Based Argumentation, with and without probabilities. *Artif. Intell.*, 293: 103449.

Cyras, K.; Oliveira, T.; Karamlou, A.; and Toni, F. 2021a. Assumption-based argumentation with preferences and goals for patient-centric reasoning with interacting clinical guidelines. *Argument Comput.*, 12(2): 149–189.

Cyras, K.; Rago, A.; Albini, E.; Baroni, P.; and Toni, F. 2021b. Argumentative XAI: A Survey. In *Proc. IJCAI*, 4392–4399. ijcai.org.

Cyras, K.; Schulz, C.; and Toni, F. 2017. Capturing Bipolar Argumentation in Non-flat Assumption-Based Argumentation. In *Proc. PRIMA*, volume 10621 of *Lecture Notes in Computer Science*, 386–402. Springer.

Cyras, K.; and Toni, F. 2016. ABA+: Assumption-Based Argumentation with Preferences. In *Proc. KR*, 553–556. AAAI Press.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2): 321–357.

Dung, P. M.; and Thang, P. M. 2010. Towards (Probabilistic) Argumentation for Jury-based Dispute Resolution. In *Proc. COMMA*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, 171–182. IOS Press.

Dung, P. M.; Thang, P. M.; and Hung, N. D. 2010. Modular argumentation for modelling legal doctrines of performance relief. *Argument Comput.*, 1(1): 47–69.

Dvorák, W.; Rapberger, A.; and Woltran, S. 2020. Argumentation Semantics under a Claim-centric View: Properties, Expressiveness and Relation to SETAFs. In *Proc. KR*, 341–350.

Dvorák, W.; and Woltran, S. 2020. Complexity of abstract argumentation under a claim-centric view. *Artif. Intell.*, 285: 103290.

Fan, X.; Liu, S.; Zhang, H.; Leung, C.; and Miao, C. 2016. Explained Activity Recognition with Computational Assumption-Based Argumentation. In *Proc. ECAI*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, 1590–1591. IOS Press.

Fan, X.; and Toni, F. 2014. A general framework for sound assumption-based argumentation dialogues. *Artif. Intell.*, 216: 20–54.

García, A. J.; and Simari, G. R. 2004. Defeasible Logic Programming: An Argumentative Approach. *Theory Pract. Log. Program.*, 4(1-2): 95–138.

Gargouri, A.; Konieczny, S.; Marquis, P.; and Vesic, S. 2021. On a notion of monotonic support for bipolar argumentation frameworks. In *Proc. AAMAS*.

Gonzalez, M. G. E.; Budán, M. C.; Simari, G. I.; and Simari, G. R. 2021. Labeled bipolar argumentation frameworks. *J. Art. Intell. Res.*, 70: 1557–1636.

Karacapilidis, N.; and Papadias, D. 2001. Computer supported argumentation and collaborative decision making: the HERMES system. *Information systems*, 26(4): 259–277.

Lehtonen, T.; Rapberger, A.; Ulbricht, M.; and Wallner, J. P. 2023. Argumentation Frameworks Induced by Assumption-based Argumentation: Relating Size and Complexity. In *Proc. KR*, 440–450.

Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2022. Algorithms for Reasoning in a Default Logic Instantiation of Assumption-Based Argumentation. In *Proc. COMMA*, volume 353 of *Frontiers in Artificial Intelligence and Applications*, 236–247. IOS Press.

Modgil, S.; and Prakken, H. 2013. A general account of argumentation with preferences. *Artif. Intell.*, 195: 361–397.

Nouioua, F.; and Risch, V. 2010. Bipolar argumentation frameworks with specialized supports. In *Proc. ICTAI*, volume 1, 215–218. IEEE.

Oren, N.; and Norman, T. J. 2008. Semantics for Evidence-Based Argumentation. In *Proc. COMMA*, volume 172 of *Frontiers in Artificial Intelligence and Applications*, 276–284. IOS Press.

Potyka, N. 2020. Bipolar abstract argumentation with dual attacks and supports. In *Proc. KR*, volume 17, 677–686.

Rapberger, A. 2020. Defining Argumentation Semantics under a Claim-centric View. In *Proc. STAIRS'20*, volume 2655 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Rapberger, A.; and Ulbricht, M. 2023. On Dynamics in Structured Argumentation Formalisms. *J. Artif. Intell. Res.*, 77: 563–643.

Toni, F. 2013. A generalised framework for dispute derivations in assumption-based argumentation. *Artif. Intell.*, 195: 1–43.

Toni, F. 2014. A tutorial on assumption-based argumentation. *Argument Comput.*, 5(1): 89–117.