Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Capturing Non-flat Assumption-based Argumentation with Bipolar SETAFs

Matti Berthold[1] , **Anna Rapberger**[2] and **Markus Ulbricht**[1]
[1]ScaDS.AI Dresden/Leipzig, Universität Leipzig, Germany
[2]Imperial College London, United Kindgom
{berthold, mulbricht}@informatik.uni-leizpig.de, a.rapberger@imperial.ac.uk

Abstract

While the flat fragment of assumption-based argumentation (ABA) is widely studied in the literature, the general, non-flat case has mostly been neglected so far. Until recently, there was no possible way to instantiate non-flat ABA in terms of an abstract argumentation framework. While this gap has been closed for complete-based ABA semantics, capturing admissible-based semantics cannot yet be achieved by looking at the relation between the instantiated arguments only; it requires augmenting arguments with their premises, hence being a semi-abstract instantiaiton. In this paper, we provide a compact and fully abstract instantiation by making use of both collective attack and support relations. Then, inspired by fundamental properties of abstract formalisms, we identify flaws of native ABA semantics in the non-flat case and provide refinements thereof, utilizing our novel instatiation.

1 Introduction

Formal argumentation is a central research area within the field of knowledge representation and reasoning. Among the most prominent argumentation formalisms is *assumption-based argumentation (ABA)*. A common restriction here is that the ABA framework (ABAF) under consideration is *flat*, i.e. assumptions cannot be entailed, only assumed to be true or not. Flat ABAFs have been studied extensively, including the study of formal properties (Cyras and Toni 2016; Dung, Kowalski, and Toni 2006; König, Rapberger, and Ulbricht 2022; Rapberger and Ulbricht 2022), computational aspects (Dvorák and Dunne 2018; Berthold, Rapberger, and Ulbricht 2023a; Berthold, Rapberger, and Ulbricht 2023b), implementations (Lehtonen, Wallner, and Järvisalo 2021a; Lehtonen, Wallner, and Järvisalo 2021b; Lehtonen et al. 2023), and applications in e.g. recommendations for patient treatment (Cyras et al. 2021) or planning (Fan 2018).

In contrast, the general non-flat fragment has only received limited attention within the last years. A severe drawback of non-flat ABAFs is that they cannot be captured by Dung's *abstract argumentation frameworks (AFs)* (Dung 1995) and consequently, the impressive body of research conducted for AFs cannot be applied to non-flat ABA (Gabbay et al. 2021). Closing this gap and finding a pleasant way to translate non-flat ABAFs into an abstract network is thus a promising approach to push forward research on general ABAFs. First steps in this direction have been done already. In (Ulbricht et al. 2024), two possible ways to instantiate non-flat ABAFs has been considered: The first one is based on bipolar AFs (BAFs) (Cayrol and Lagasquie-Schiex 2005), but only works for complete-based semantics. A significant advantage of this approach is, however, the lower computational complexity in BAFs compared to non-flat ABA (Cyras, Heinrich, and Toni 2021). Consequently, after the computational cost of constructing the BAF has been paid, the reasoning problems can be solved more efficiently. The second approach extends BAFs to so-called premise-augmented BAFs (pBAFs). While pBAFs capture non-flat ABA under all standard semantics, the drawback is that in pBAFs, we require more technical baggage.

In this work, we refine the instantiation of non-flat ABA via an abstract formalism that captures non-flat ABA under all standard semantics, while maintaining the low computational cost of Dung's AFs. To this end we propose *bipolar SETAFs* (BSAFs), an extension of SETAFs (Nielsen and Parsons 2006; Bikakis et al. 2021) that is also capable of modeling collective support among arguments. Studying the formal properties of BSAFs we will realize that they inherit several undesired properties from ABA. We demonstrate how to refine the BSAF semantics in a natural way, avoiding the aforementioned drawbacks. Having achieved this, we discuss conceivable consequences of this towards refining the non-flat ABA semantics.

2 Background

We recall the technical definitions of (ABA) (Čyras et al. 2018). We assume a *deductive system*, i.e. a tuple $(\mathcal{L}, \mathcal{R})$, where $\mathcal{L}$ is a set of atoms and $\mathcal{R}$ is a set of inference rules over $\mathcal{L}$. A rule $r \in \mathcal{R}$ has the form $a_0 \leftarrow a_1, \ldots, a_n$, s.t. $a_i \in \mathcal{L}$ for all $0 \leq i \leq n$; $head(r) := a_0$ is the *head* and $body(r) := \{a_1, \ldots, a_n\}$ is the (possibly empty) *body* of r.

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

In this work we focus on frameworks which are *finite*, i.e. for which $\mathcal{L}$ and $\mathcal{R}$ are finite. An atom $p \in \mathcal{L}$ is tree-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 s.t. i) the root of t is labeled with p, ii) the set of labels for the leaves of t is equal to S or $S \cup \{\top\}$, and iii) for each node

v that is not a leaf of t there is a rule $r \in R$ such that v is labeled with $head(r)$ and labels of the children correspond to $body(r)$ or $\top$ if $body(r) = \emptyset$. We write $S \vdash p$ iff there exists $R \subseteq \mathcal{R}$ such that $S \vdash_R p$.

By $Th_{\mathcal{D}}(S) := \{p \in L \mid \exists S' \subseteq S : S' \vdash p\}$ we denote the set of all conclusions derivable from an assumption-set S in an ABAF $\mathcal{D}$. A set $S \subseteq \mathcal{A}$ *attacks* a set $T \subseteq \mathcal{A}$ if there are $S' \subseteq S$ and $a \in T$ s.t. $S' \vdash \overline{a}$; if S attacks $\{a\}$ we say S attacks a. Set S is conflict-free ($S \in cf(\mathcal{D})$) if it does not attack itself. Given $S \subseteq \mathcal{A}$, the closure $cl(S)$ of S is $cl(S) := Th_{\mathcal{D}}(S) \cap \mathcal{A}$. We write $cl(a)$ instead of $cl(\{a\})$ for singletons. We call $S \subseteq \mathcal{A}$ *closed* if $S = cl(S)$.

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.

Definition 2.2. *A set S of assumptions defends an assumption a iff for each closed set T which attacks a, we have S attacks T; S defends itself iff S defends each $b \in S$.*

A set E of assumptions is *admissible* ($E \in ad(F)$) iff E is conflict-free, closed and defends itself. We next recall grounded, complete, preferred, and stable ABA semantics.

Definition 2.3. *Let $\mathcal{D}$ be an ABAF and let $S \in ad(\mathcal{D})$. Then*

- *$S \in co(\mathcal{D})$ iff it contains every assumption it defends;*
- *$S \in gr(\mathcal{D})$ iff S is $\subseteq$-minimal in $co(\mathcal{D})$;*
- *$S \in pr(\mathcal{D})$ iff S is $\subseteq$-maximal in $ad(\mathcal{D})$;*
- *$S \in stb(\mathcal{D})$ iff S is closed and attacks each $x \in \mathcal{A}(\mathcal{D}) \setminus S$.*

Example 2.4. *We consider an ABAF $\mathcal{D}$ with literals $\mathcal{L} = \{a, b, c, d, e, \overline{a}, \overline{b}, \overline{c}, \overline{d}, \overline{e}\}$, assumptions $\mathcal{A} = \{a, b, c, d, e\}$, their contraries $\overline{a}$, $\overline{b}$, $\overline{c}$, $\overline{d}$ and $\overline{e}$, respectively, and rules*

$$\overline{e} \leftarrow e \qquad e \leftarrow a, b \qquad \overline{a} \leftarrow d \qquad \overline{c} \leftarrow d$$
$$\overline{e} \leftarrow b, c \qquad \overline{d} \leftarrow a \qquad \overline{d} \leftarrow c$$

The set $\{a\}$ is admissible as it defends itself against $\{d\}$; it is not complete, however, since it also defends c and b. The framework has two complete extension $\{b\}$ and $\{b, d\}$. As a result $\{b\}$ is the grounded extension. Moreover, the framework has three preferred extensions $\{a, c\}$, $\{b, c\}$ and $\{b, d\}$, and no stable extensions.

3 Bipolar SETAF

In this section we introduce Bipolar SETAF as a faithful abstraction of bipolar AFs as used by (Ulbricht et al. 2024), before arguing that they capture non-flat ABAFs.

BSAFs combine the ideas underlying SETAFs and BAFs in the following sense: Instead of only considering an attack relation, there is also a notion of support. Moreover, BSAFs can model *collective* attacks and supports.

Definition 3.1. *A bipolar set-argumentation framework (BSAF) is a tuple $F = (A, R, S)$, where A is a finite set of arguments, $R \subseteq 2^A \times A$ is the attack relation and $S \subseteq 2^A \times A$ is the support relation.*

Given a BSAF $F = (A, R, S)$, we call $\mathcal{F} := (A, R)$ the underlying SETAF of F. If $|T| = 1$ holds for each $(T, h) \in R \cup S$, then we naturally identify F with a biploar AF (BAF). As usual, we say E *attacks* E' in F, iff there are $T \subseteq E$ and $h \in E'$, s.t. $(T, h) \in R$; E is conflict-free ($E \in cf(F)$) if it does not attack itself.

Definition 3.2. *Given a BSAF $F = (A, R, S)$ and a set $E \subseteq A$ of arguments. With*

$$supp_F(E) := E \cup \{h \in A \mid \exists (T, h) \in S : T \subseteq E\}$$

we call $cl_F(E) = \bigcup_{i \geq 1} supp^i_F(E)$ the closure *of E; E is* closed *if $cl_F(E) = E$.*

Let us now head to the semantics of BSAFs. A set $E \subseteq A$ *defends* $a \in A$ if for each closed set $E' \subseteq A$ attacking a, E attacks E'; E *defends* E' if E defends each $a \in E'$. We define the characteristic function by $\Gamma_F(E) := \{a \in A \mid E$ defends a in $F\}$. If clear from the context, we omit the subscript F for Γ and cl. A conflict-free set E is admissible ($E \in ad(F)$) iff E is closed and defends itself.

Definition 3.3. *Let F be an BSAF and let $E \in ad(F)$.*

- *$E \in co(F)$ iff E contains every assumption it defends;*
- *$E \in gr(F)$ iff E is $\subseteq$-minimal in $co(F)$;*
- *$E \in pr(F)$ iff E is $\subseteq$-maximal in $ad(F)$;*
- *$E \in stb(F)$ iff E attacks each $x \in A \setminus E$.*

Graphically, we depict the attack relation of a BSAF by solid edges and the support relation by dashed edges.

Example 3.4. *Consider the BSAF $F = (A, R, S)$, where $A = \{a, b, c, d, e\}$, $R = \{(\{a\}, d), (\{d\}, a), (\{d\}, c), (\{c\}, d), (\{b, d\}, e), (\{e\}, e)\}$, and $S = \{(\{a, b\}, e)\}$.*

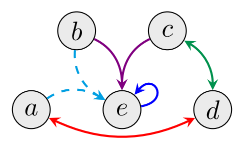

Here we have $ad(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{a, c\}, \{b, c\}, \{b, d\}\}$, $co(F) = \{\{b\}, \{b, d\}\}$, $gr(F) = \{\{b\}\}$, $pr(F) = \{\{a, c\}, \{b, c\}, \{b, d\}\}$, and $stb(F) = \emptyset$.

Our BSAFs faithfully generalize the BAFs as defined recently: if $|T| = 1$ holds for each $(T, h) \in R \cup S$, then the BSAF amounts to a BAF as defined in (Ulbricht et al. 2024).

3.1 ABA and BSAF

In the following we show that BSAFs are capable of instantiating non-flat ABA under all standard semantics. This is in contrast to the BAF semantics as considered in (Ulbricht et al. 2024) which lack the expressive power to spot the necessary collective support behavior in non-flat ABA.

Capturing non-flat ABAFs by means of BSAFs is quite natural, as shown in the next definition.

Definition 3.5. *Let $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF. Then we set $F_{\mathcal{D}} := (A, R, S)$, where $A := \mathcal{A}$ and*

$$R := \{(T, h) \mid h \in \mathcal{A},\ T \vdash \overline{h}\},$$
$$S := \{(T, h) \mid h \in \mathcal{A},\ T \vdash h\}\}$$

Example 3.6. *Recall the ABAF given in Example 2.4. It is instantiated by the BSAF in Example 3.4.*

The next theorem states that the BSAF $F_{\mathcal{D}}$ indeed captures reasoning in $\mathcal{D}$ in the expected way.

Theorem 3.7. *Given a semantics $\sigma \in \Sigma$, and an ABAF $\mathcal{D}$. Then $\sigma(\mathcal{D}) = \sigma(F_{\mathcal{D}})$.*

Vice versa, each BSAF F can be captured by some ABAF $\mathcal{D}_F$ as follows.

Definition 3.8. *Let $F = (A, R, S)$ be a BSAF, then we define $\mathcal{D}_F := (\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ })$, where $\mathcal{A} := A$ and*

$$\begin{aligned}\mathcal{L} &:= A \cup \{\overline{a} \mid a \in A\},\\ \mathcal{R} &:= \{h \leftarrow T \mid (T, h) \in S, h \in \mathcal{A}\} \cup\\ &\quad\ \{\overline{h} \leftarrow T \mid (T, h) \in R, h \in \mathcal{A}\},\\ \bar{\ } &:= \{(a, \overline{a}) \mid a \in \mathcal{A}\}.\end{aligned}$$

Example 3.9. *Recall the BSAF in Example 3.4. Using the construction above, it can be translated into the ABAF given in Example 2.4 (up to possible renaming of the newly introduced contraries).*

Theorem 3.10. *Given a semantics $\sigma \in \Sigma$, and a BSAF F. Then $\sigma(F) = \sigma(\mathcal{D}_F)$.*

3.2 Properties of BSAF

Let us now discuss formal properties of BSAFs. Regarding basic semantics relations, we observe that admissible, preferred, and stable semantics are related to each other as we know it from traditional AF semantics.

Proposition 3.11. *Let F be any BSAF. Then it holds that $stb(F) \subseteq pr(F) \subseteq ad(F) \subseteq cf(F)$.*

Another important observation about BSAFs is the nature of defense. Recall that defense is necessary only against closed sets of arguments. Due to the potentially involved collective support relation, it might seem hard at first glance to search for all such closed attacking sets. However, there is a simple characterization, given as follows.

Lemma 3.12. *Let $F = (A, R, S)$ be a BSAF and let $E \subseteq A$. Then $a \in \Gamma(E)$ iff for each $(T, a) \in R$, E attacks $cl(T)$.*

From this we can deduce that whether or not some set E defends itself can be verified in polynomial time. Consequently, most reasoning problems are not harder compared to Dung AFs (and thus easier that in a non-flat ABAF). The only exception is gr since we cannot compute it by iterating the characteristic function. Hence, we have to check whether the input set is minimal in $co(F)$ from scratch.

Theorem 3.13. *Deciding whether*

- *$E \in \sigma(D)$ is tractable for $\sigma \in \{ad, co, stb\}$, coNP-complete for pr, and $\mathsf{D}_1^{\mathsf{P}}$-complete for gr;*
- *a is credulously accepted is NP-complete for $\sigma \in \{ad, co, pr, stb\}$ and $\mathsf{D}_1^{\mathsf{P}}$-complete for gr;*
- *a is skeptically accepted is coNP-complete for ad, $\mathsf{D}_1^{\mathsf{P}}$-complete for $\sigma \in \{co, gr, stb\}$ and Π_P^2-complete for pr.*

However, from the semantics correspondence between BSAFs and ABA, some undesired properties of BSAFs are also immediate. We observe admissible extensions may not always exist, as the following example shows.

Example 3.14. *Let F be a BSAF with $(\emptyset, a) \in S$ and $(b, a), (b, b) \in R$. Not even the empty set can be accepted admissibly: $\emptyset$ supports a which is attacked by the (otherwise undisputed) self-attacker b.*

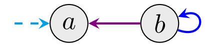

Moreover, even if admissible extensions exist, it can be the case that BSAFs may have no complete (thus also no grounded) extensions.

Example 3.15. *Consider a BSAF F with supports $(\{a, b\}, d), (\{a, c\}, d) \in S$ and attack $(\{d\}, d) \in R$.*

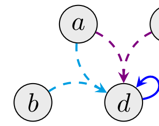

The admissible sets are $\emptyset$, $\{a\}$, $\{b\}$, $\{c\}$, $\{b, c\}$. However, neither a, b, nor c are attacked; hence, all arguments are defended by $\emptyset$. However, the set $\{a, b, c\}$ supports the self-attacker d. Consequently, F has no complete extensions.

Overall, we observe the following (undesired) properties:

- It might happen that $ad(F) = \emptyset$ and $pr(F) = \emptyset$.
- It might happen that $\emptyset \notin ad(F)$.
- Even if $ad(F) \neq \emptyset$, it might happen that $co(F) = \emptyset$.
- Even if $co(F) \neq \emptyset$, it might happen that $|gr(F)| > 1$.
- Even if $co(F) \neq \emptyset$, it might happen that some $E \in pr(F)$ is not complete.

One of the main technical ingredients to ensure the beneficial behavior of the classical abstract argumentation semantics is Dung's fundamental lemma (Dung 1995). In a nutshell, it amounts to the following lemma.

Lemma 3.16. *(Dung 1995) Let $F = (A, R)$ be an AF and let $E \subseteq A$ and $a \in A$. If $E \in ad(F)$ and E defends a, then $E \cup \{a\} \in ad(F)$.*

Indeed, for non-flat ABA (as well as our BSAF semantics), the fundamental lemma does not hold.

4 Alternative BSAF Semantics

In the last section we observed that BSAFs inherit some undesirable properties from ABA. Intuitively, the reason for the undesired behavior is that the notions of *closure* and *defense* might induce conflicting requirements. In this section, we set out to fix (most of) these shortcomings by better coordinating these two concepts. We first propose a stricter notion for closure; however, by doing so, we introduce undesired behavior already for admissible-based semantics, as we discuss. We then propose instead an appropriate adjustment of the closure, which results in a family of novel semantics for BSAFs and, by our translation, also for ABA.

4.1 Fixing Admissible Semantics

Let us first focus on the notion of support. As observed in Example 3.14, it can be the case that no admissible extension exists, e.g., if the empty set supports an argument that is impossible to defend. We therefore stipulate that an argument a only counts as supported by a set E if the supporting set is strong enough to defend a against each attack, and propose a refined closure notion which takes this into account.

Definition 4.1. *Given a BSAF $F = (A, R, S)$, $E \subseteq A$, and $a \in A$. Then E Γ-supports a iff $a \in cl(E)$ and $a \in \Gamma(E)$; E is Γ-closed iff E contains all arguments it Γ-supports.*

For a BSAF F, we call a set $E \subseteq A$ of arguments Γ-admissible ($E \in ad_\Gamma(F)$) iff E defends itself and is Γ-closed. The new Γ-support notion naturally induces the refined complete-based semantics co_Γ, gr_Γ and pr_Γ.

Example 4.2. *Heading back to the BSAF from Example 3.14, we observe that the empty set is now acceptable: since a cannot be defended against b we do not include it in the Γ-closure of $\emptyset$. Hence $\emptyset$ is admissible.*

Our adjusted semantics is successful in ensuring the existence of at least one admissible extension.

Proposition 4.3. *Let F be a BSAF. Then $ad_\Gamma(F) \neq \emptyset$ as well as $pr_\Gamma(F) \neq \emptyset$.*

We therefore successfully adjusted BSAF-semantics such that there are always admissible (and hence also preferred) extensions. From a theoretical perspective this would be a solid foundation to also consider adjustments to the completeness notion. However, a closer inspection reveals certain obstacles of our novel semantics. Revisiting Example 3.15 indicates that the proposed fix is too liberal; each subset of $\{a, b, c\}$ is accepted under Γ-admissible semantics.

Example 4.4. *Consider again the BSAF from Example 3.15. The Γ-admissible sets are $\emptyset$, $\{a\}$, $\{b\}$, $\{c\}$, $\{b, c\}$, but also, $\{a, b\}$, $\{a, c\}$, and $\{a, b, c\}$. The collective support towards d can be ignored since d cannot be defended.*

In the above example, contrary to any intuition, the arguments a, b and c may be jointly accepted, even without d, even though they pose a sufficient condition for d to hold.

Moreover, although our new support notion requires defense, it holds that $ad_\Gamma(F) \neq \emptyset$ does not imply $co_\Gamma(F) \neq \emptyset$, the following example illustrates.

Example 4.5. *Let F be a BSAF with supports $(\{b\}, d)$, $(\{a\}, d) \in S$ and attacks $(\{a\}, b)$, $(\{a\}, b)$, $(\{c\}, d) \in R$.*

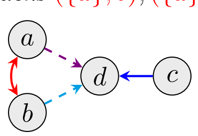

Here we have $ad(F) = ad_\Gamma(F) = \{\emptyset, \{c\}\}$. On the other hand, $co(F) = co_\Gamma(F) = \emptyset$: c (and its closure) is unattacked, hence it is in any complete extension. Also, c defends b from a since it attacks the closure of a. Vice versa c defends a from b, as well. Then the closure of $\{a, b, c\}$ is not conflict-free; hence, no complete extension exists.

The example indicates that the original notion of defense might be too liberal to confirm our intuition. In the following we therefore propose an adjustment which interleaves support and defense but this time by proposing a novel defense notion while sticking to the original closure definition.

4.2 Fixing Complete Semantics

In the previous section we successfully redefined the support s.t. admissible semantics are universally defined. However, we were still facing unwanted behavior in the semantics.

Let us now adjust the notion of defense in order to circumvent the aforementioned issues. Recall that in BSAFs, complete extensions do not necessarily exist, even if $ad(F) \neq \emptyset$. We observe that the absence of complete extensions in such a situation stems from the fact that checking $a \in \Gamma(E)$ ignores the consequences of including a into our given extension E. More specifically, the closure $cl(E \cup \{a\})$ might contain arguments we did not yet take into consideration. We thus refine defense by stipulating that $cl(E \cup \{a\}) \subseteq \Gamma(E)$ in order for E to defend a.

Definition 4.6. *Given a BSAF $F = (A, R, S)$, $E \subseteq A$, and $a \in A$. Then E Δ-defends a if $cl(E \cup \{a\}) \subseteq \Gamma(E)$.*

We note that each admissible (and preferred) extension Δ-defends itself. The new Δ-defense notion naturally induces the refined complete-based semantics co_Δ and gr_Δ.

Definition 4.7. *Let F be an BSAF and let $E \in ad(F)$.*

- *$E \in co_\Delta(F)$ iff it contains each Δ-defended argument;*
- *$E \in gr_\Delta(F)$ iff E is $\subseteq$-minimal in $co_\Delta(F)$.*

Let us first reconsider the BSAF from Example 3.15.

Example 4.8. *Recall the BSAF F from Example 3.15. The set $S_1 = \{a\}$ does not Δ-defend any further arguments: although b is defended (because it is unattacked) we have that $\{a, b\}$ supports d; thus $cl(S_1 \cup \{b\}) = \{a, b, d\} \not\subseteq \Gamma(S_1) = \{a, b, c\}$. The analogous observation holds for c. Therefore, we have that $co_\Delta(F) = gr_\Delta(F) = \{\{a\}, \{b, c\}\}$.*

Let us next revisit Example 4.5.

Example 4.9. *Let us consider again the BSAF from Example 4.5. In the previous case, c defended both a and b because their closure $cl(a)$ and $cl(b)$ were attacked. With our new adjustment, the argument c does not defend the arguments anymore: by definition, $\{c\}$ defends a if $cl(\{c, a\})$ is defended by $\{c\}$. The closure of $\{c, a\}$ is the set $\{a, c, d\}$ which is attacked by c. Therefore, the set $\{c\}$ is already Δ-complete. We have $co_\Delta(F) = \{\{c\}\}$.*

Overall, we observe that the proposed Δ-defense notion guarantees that each admissible set has a completion; moreover, a restricted version of the fundamental lemma holds, and each preferred set is complete.

Proposition 4.10. *Let F be a BSAF.*

- *If $ad(F) \neq \emptyset$ then $co_\Delta(F) \neq \emptyset$.*
- *If E is admissible and Δ-defends $a \in A$, then there is some admissible set $E' \in ad(F)$ s.t. $E' \supseteq E \cup \{a\}$.*
- *It holds that $pr(F) \subseteq co_\Delta(F) \subseteq ad(F)$.*

These results demonstrate that the newly introduced semantics admit an intuitive behavior whilst confirming to natural adjustments of the defense mechanism.

As a final remark regarding the novel semantics defined in Sections 4.1 and 4.2 we want to mention that the adjusted versions of support and defense merely introduce an additional polynomial check when computing extensions. Therefore, we strongly conjecture that the computational complexity persists; however, a thorough investigation of this issue is left for future work.

4.3 Consequences for ABA

Due to the close relation between BSAFs and ABAFs that we have established in the previous section, we can directly translate our newly introduced semantics and our results to the realm of ABA. Thereby, we provide an adjustment of complete(-based) semantics that adheres to desired behavior of argumentation semantics.

Definition 4.11. *Given an ABAF* $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$*,* $E \subseteq \mathcal{A}$*, and* $a \in \mathcal{A}$*. Then* E Δ*-defends* a *if* $cl(E \cup \{a\}) \subseteq \Gamma(E)$*.*

As it is the case for BSAFs, it holds that each admissible (and preferred) assumption extension Δ-defends itself. We are ready to define the refined complete-based semantics co_Δ and gr_Δ for ABA.

Definition 4.12. *Let* $\mathcal{D}$ *be an ABAF and let* $E \in ad(\mathcal{D})$*.*

- $E \in co_\Delta(\mathcal{D})$ *iff it contains each* Δ*-defended assumption;*
- $E \in gr_\Delta(\mathcal{D})$ *iff* E *is* $\subseteq$*-minimal in* $co_\Delta(\mathcal{D})$*.*

The direct correspondence between (non-flat) ABA and BSAF holds under our newly defined semantics. As a consequence from the results established in the previous section, our novel semantics satisfy the following properties.

Proposition 4.13. *Let* $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^{-})$ *be an ABAF.*

- *If* $ad(\mathcal{D}) \neq \emptyset$ *then* $co_\Delta(\mathcal{D}) \neq \emptyset$*.*
- *If* E *is admissible and* Δ*-defends* $a \in \mathcal{A}$*, then there is some admissible set* $E' \in ad(\mathcal{D})$ *s.t.* $E' \supseteq E \cup \{a\}$*.*
- *It holds that* $pr(\mathcal{D}) \subseteq co_\Delta(\mathcal{D}) \subseteq ad(\mathcal{D})$*.*

5 Conclusion

We presented a compact and fully abstract expansion of Dung AFs that capture general (non-flat) assumption-based argumentation frameworks. In order to overcome certain flaws of the traditional ABA semantics, we discussed alternative ways of connecting defense and support of arguments. We identified an intuitive adjustment of the defense notion (Δ-*defense*) and showed that the induced complete-based semantics admit intuitive properties.

In the literature, there is a fruitful debate as to how define BAF semantics suitably (Karacapilidis and Papadias 2001; Cayrol and Lagasquie-Schiex 2005; Amgoud et al. 2008; Oren and Norman 2008; Nouioua and Risch 2011; Cayrol and Lagasquie-Schiex 2020); however, since our aim was to capture non-flat ABAFs, our definition of the semantics was inspired by those. Nonetheless, an interesting future work direction consists in comparing our semantics to the existing ones. Evidence-based argumentation frameworks (Oren and Norman 2008) are syntactically similar to BSAFs in the sense that they also features set-to-argument supports and attacks. The support relation there, however, works differently, as an argument is accepted only if it is supported through a chain of arguments tracking back to some ground truth; i.e., their paper considers necessary support whereas our notions are closer related to deductive support. It would be interesting to find a formalism that encompasses both of these support interpretations. Another interesting avenue for future research would be to see how BSAFs behave with respect to the recently conceived principles of bipolar argumentation (Yu et al. 2023).

Acknowledgements

This research has been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101020934), 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.

References

Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems*.

Berthold, M.; Rapberger, A.; and Ulbricht, M. 2023a. Forgetting aspects in assumption-based argumentation. In *Proceedings (KR-23)*.

Berthold, M.; Rapberger, A.; and Ulbricht, M. 2023b. On the expressive power of assumption-based argumentation. In *Proceedings of (JELIA-23)*.

Bikakis, A.; Cohen, A.; Dvorák, W.; Flouris, G.; and Parsons, S. 2021. Joint attacks and accrual in argumentation frameworks. *IfCoLog Journal of Logics and their Applications*.

Bondarenko, A.; Dung, P. M.; Kowalski, R. A.; and Toni, F. 1997. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*.

Cayrol, C., and Lagasquie-Schiex, M. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Proceedings of (ECSQARU-05)*.

Cayrol, C., and Lagasquie-Schiex, M. 2020. Logical encoding of argumentation frameworks with higher-order attacks and evidential supports. *International Journal on Artificial Intelligence Tools*.

Cyras, K., and Toni, F. 2016. ABA+: assumption-based argumentation with preferences. In *Proceedings of (KR-16)*.

Čyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2018. Assumption-based argumentation: Disputes, explanations, preferences. In *Handbook of Formal Argumentation*, volume 1. chapter 7.

Cyras, K.; Oliveira, T.; Karamlou, A.; and Toni, F. 2021. Assumption-based argumentation with preferences and goals for patient-centric reasoning with interacting clinical guidelines. *Argument & Computation*.

Cyras, K.; Heinrich, Q.; and Toni, F. 2021. Computational complexity of flat and generic assumption-based argumentation, with and without probabilities. *Artificial Intelligence*.

Dung, P. M.; Kowalski, R. A.; and Toni, F. 2006. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*.

Dvorák, W., and Dunne, P. E. 2018. Computational problems in formal argumentation and their complexity. In *Handbook of Formal Argumentation*. also appears in IfCoLog Journal of Logics and their Applications.

Fan, X. 2018. A temporal planning example with assumption-based argumentation. In *Proceedings of (PRIMA-18)*.

Gabbay, D.; Giacomin, M.; Simari, G. R.; and Thimm, M., eds. 2021. *Handbook of Formal Argumentation*, volume 2.

Karacapilidis, N., and Papadias, D. 2001. Computer supported argumentation and collaborative decision making: the HERMES system. *Information Systems*.

König, M.; Rapberger, A.; and Ulbricht, M. 2022. Just a matter of perspective. In *Proceedings of (COMMA-22)*.

Lehtonen, T.; Rapberger, A.; Ulbricht, M.; and Wallner, J. P. 2023. Argumentation frameworks induced by assumption-based argumentation: Relating size and complexity. In *Proceedings of (KR-23)*.

Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2021a. Declarative algorithms and complexity results for assumption-based argumentation. *Journal of Artificial Intelligence Research*.

Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2021b. Harnessing incremental answer set solving for reasoning in assumption-based argumentation. *Theory and Practice of Logic Programming*.

Nielsen, S. H., and Parsons, S. 2006. A generalization of dung's abstract framework for argumentation: Arguing with sets of attacking arguments. In *Proceedings of (ArgMAS-06)*.

Nouioua, F., and Risch, V. 2011. Argumentation frameworks with necessities. In *Proceedings of (SUM-11)*.

Oren, N., and Norman, T. J. 2008. Semantics for evidence-based argumentation. In *Proceedings of (COMMA-08)*.

Rapberger, A., and Ulbricht, M. 2022. On dynamics in structured argumentation formalisms. In *Proceedings of (KR-22)*.

Ulbricht, M.; Potyka, N.; Rapberger, A.; and Toni, F. 2024. Non-flat ABA is an instance of bipolar argumentation. In *Proceedings of (AAAI-24)*.

Yu, L.; Anaissy, C. A.; Vesic, S.; Li, X.; and van der Torre, L. 2023. A principle-based analysis of bipolar argumentation semantics. In *Proceedings of (JELIA-23)*.