

Redefining ABA^+ Semantics via Abstract Set-to-Set Attacks

Yannis Dimopoulos¹, Wolfgang Dvořák², Matthias König², Anna Rapberger³, Markus Ulbricht⁴,
Stefan Woltran²

¹University of Cyprus, Department of Computer Science

²TU Wien, Institute of Logic and Computation

³Imperial College London, Department of Computing

⁴Leipzig University, Department of Computer Science

yannis@cs.ucy.ac.cy, mulbricht@informatik.uni-leipzig.de, a.rapberger@imperial.ac.uk
{wolfgang.dvorak, matthias.koenig, stefan.woltran}@tuwien.ac.at

Abstract

Assumption-based argumentation (ABA) is a powerful defeasible reasoning formalism which is based on the interplay of assumptions, their contraries, and inference rules. ABA with preferences (ABA^+) generalizes the basic model by allowing a qualitative comparison of assumptions. The integration of preferences however comes with a cost. In ABA^+ , the evaluation under two central and well-established semantics—grounded and complete semantics—is not guaranteed to yield an outcome. Moreover, while ABA frameworks without preferences allow for a graph-based representation in Dung-style frameworks, an according instantiation for general ABA^+ frameworks has not been established so far. In this work, we tackle both issues: First, we develop a novel abstract argumentation formalism based on set-to-set attacks. We show that our so-called Hyper Argumentation Frameworks (HYPAFs) capture ABA^+ . Second, we propose relaxed variants of complete and grounded semantics for HYPAFs that yield an extension for all frameworks by design, while still faithfully generalizing the established semantics of Dung-style Argumentation Frameworks. We exploit the newly established correspondence between ABA^+ and HYPAFs to obtain variants for grounded and complete ABA^+ semantics that are guaranteed to yield an outcome. Finally, we discuss basic properties and provide a complexity analysis. Along the way, we settle the computational complexity of several ABA^+ semantics.

1 Introduction

Formal argumentation is a major research area in knowledge representation and reasoning, with applications in various fields in the realm of Artificial Intelligence (Bench-Capon, Prakken, and Sartor 2009; Baroni et al. 2018). The close interplay between rule-based systems and graph-based methods is thereby key to exploit the full potential of argumentative methods. Rule-based formalisms are crucial to understand and evaluate complex dependencies between defeasible elements of a knowledge base. *Assumption-Based Argumentation (ABA)* (Cyras et al. 2018) is one of the leading rule-based argumentation formalisms. Key elements are assumptions, their contraries, and inference rules; they form the building blocks to construct arguments and to identify conflicts in the given knowledge base. Acceptability of the

assumptions is evaluated by employing so-called *argumentation semantics*. Graph-based representations, on the other hand, offer an intuitive and easy-to-understand model of argumentative scenarios. One of the most popular models are *abstract Argumentation Frameworks (AFs)* (Dung 1995) which are directed graphs where nodes correspond to arguments and the arcs encode attacks (directed conflicts) between them. In recent years, researchers have proposed and further developed numerous generalizations. A notable one are AFs with *collective attacks (SETAFs)* due to Nielsen and Parsons (2006). SETAFs allow for situations where sets of arguments jointly attack a single argument.

The *instantiation* of knowledge bases in terms of abstract frameworks yields access to the visual advantages of graph-based models and allows the application of well-developed algorithms, tailored to abstract frameworks.

Example 1. *Alice wants to go on vacation to Tokyo. Ideally, she would like to visit for two weeks and stay in a 5-star hotel; however, such a long stay in a luxurious hotel is too expensive for her. We can formalize this as an ABA framework with assumptions t (“Alice travels to Tokyo”), w (“Alice visits for two weeks”), and h (“Alice stays in a 5-star hotel”) and their contraries \bar{t} , \bar{w} , and \bar{h} , expressing the negation of the respective assumptions. The rule $r : (\bar{t} \leftarrow w, h)$ formalizes that two weeks in a 5-star hotel implies that Alice cannot travel to Tokyo. An intuitive representation of conflicts in this framework can be obtained via the following simple SETAF (König, Rapberger, and Ulbricht 2022):*




The arrow from the set containing w and h to t indicates the joint attack. The set $\{w, h\}$ can be accepted since they are unattacked whereas t is defeated by them. Hence, Alice won't be able to go to Tokyo, given her trip specifications.

Suppose we also want to model different preferences of Alice, e.g. she might think that it would be better to have a short stay in Tokyo than none at all, i.e., she prefers going to Tokyo over staying for two weeks. There is a variety of approaches to handle preferences in different argumentation formalisms (Modgil and Prakken 2014; Besnard and Hunter 2014; Amgoud and Vesic 2014). *ABA with preferences (ABA^+)* (Cyras and Toni 2016) generalizes ABA

by allowing qualitative comparisons between assumptions. ABA^+ has been studied extensively (Cyras et al. 2018) and computational analyses as well as efficient solvers are available (Lehtonen, Wallner, and Järvisalo 2021a,b).

Example 2. *In our case, we get the preference ordering $t > w$. Preferences in ABA^+ are handled as follows: Since t is preferred over w , the attack from $\{w, h\}$ to t is reversed. The underlying rationale is that, on the one hand, the conflict between t and $\{w, h\}$ must be preserved; however, on the other hand, Alice prefers t over w and thus the reversed attack better captures the intended meaning.*



However, this situation cannot be expressed as SETAF since the head of the reversed attack contains multiple elements, which is beyond the scope of the traditional model.

While the connection between ABA and abstract graph-based representations is well understood (Caminada et al. 2015; Cyras et al. 2018; König, Rapberger, and Ulbricht 2022), an instantiation procedure for ABA^+ is not yet developed. Bao, Cyras, and Toni (2017) discuss an instantiation for a restricted class of ABA^+ frameworks satisfying *weak contraposition*; however, they show that their method cannot be used for ABA^+ in its full generality (e.g., Example 2 violates weak contraposition). As demonstrated in the above example, also SETAFs fail to capture ABA^+ .

There is, however, an even bigger issue: In contrast to most established argumentation formalisms, an ABA^+ framework does not necessarily admit a *complete* extension. Intuitively, a set S of assumptions is complete if i) S does not have internal conflicts, ii) each assumption $a \in S$ is *defended*, i.e., S refutes each set of assumptions that objects against a , and iii) S contains each assumption it defends. This notion is a cornerstone in argumentation theory and the absence of complete extensions is a major drawback in ABA^+ . Even the (simple) example above does not admit a complete extension: Each assumption w , h , and t is unattacked (and thus defended) individually, but the whole set $\{w, h, t\}$ has the internal conflict from t to $\{w, h\}$.

In this paper, we tackle both issues. For this, we generalize traditional AFs by allowing for set-to-set attacks. Our so-called *Hyper Argumentation Frameworks (HYPAFs)* will serve as theoretical basis to develop adaptations of complete semantics for ABA^+ that avoid the undesired behavior. Our novel *closure operator* Θ_F will be central for our goal as it adopts a more cautious approach compared to usual defense.

Our main contributions are as follows.

- We develop an abstract representation for ABA^+ . For this, we propose a generalization of AFs that allows attacks between sets. We show that HYPAFs capture ABA^+ by providing a semantics-preserving translation.
- We propose an adaptation of complete semantics for HYPAFs by refining the characteristic function; our semantics retain desirable properties of their AF counterparts.
- We exploit the novel connection between HYPAF and ABA^+ to obtain new complete ABA^+ semantics that guarantee extension existence.

- We discuss properties of the novel ABA^+ semantics and conduct a complexity analysis. Along the way, we fill some gaps in the computational complexity landscape of the ABA^+ semantics that have been open so far.

2 Preliminaries

We recall Assumption-Based Argumentation with Preferences (ABA^+) (Cyras and Toni 2016) and briefly recall a graph-based instantiation that utilizes argumentation frameworks with collective attacks (SETAFs) (Nielsen and Parsons 2006). This connection is akin to our instantiation method that we will introduce in Section 4.

ABA⁺ 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 $a_0 \leftarrow a_1, \dots, a_n$, with $a_i \in \mathcal{L}$, $head(r) = a_0$ is the head and $body(r) = \{a_1, \dots, a_n\}$ is the (possibly empty) body of r .

Definition 3. *An ABA^+ framework is a tuple $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\cdot}, \leq)$, where*

- $(\mathcal{L}, \mathcal{R})$ is a deductive system,
- $\mathcal{A} \subseteq \mathcal{L}$ a non-empty set of assumptions,
- $\bar{\cdot} : \mathcal{A} \rightarrow \mathcal{L}$ is the contrary function,
- and \leq is a reflexive, transitive binary relation on \mathcal{A} .

As usual, we write $a < b$ if $a \leq b$ and $b \not\leq a$; D is an ABA framework (without preferences) if \leq is empty.

In this work, we focus on frameworks which are *flat*, i.e., $head(r) \notin \mathcal{A}$ for all $r \in \mathcal{R}$, and *finite*, i.e., $\mathcal{L}, \mathcal{R}, \mathcal{A}$ are finite.

We say that a sentence $p \in \mathcal{L}$ 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 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 $body(r) = \emptyset$. We call $S \vdash p$ an *argument* iff there is a tree-derivation $S \vdash_R p$.

Example 4. *Let us revisit Example 1. The corresponding formalization in terms of ABA^+ is $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\cdot}, \leq)$ with $\mathcal{L} = \{t, \bar{t}, w, \bar{w}, h, \bar{h}\}$, $\mathcal{R} = \{\bar{t} \leftarrow w, h\}$, $\mathcal{A} = \{t, w, h\}$, and preference $t > w$. The contrary function $\bar{\cdot}$ is implicitly given by the names of the atoms (e.g., the contrary of t is \bar{t}).*

Inspired by Lehtonen, Wallner, and Järvisalo (2021a), we introduce ABA^+ attacks via *normal* and *reversed* attacks. ABA without preferences has only normal attacks. Preferences might reverse attacks, as formalized next in item (ii).

Definition 5. *Given an ABA^+ framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\cdot}, \leq)$, an assumption set $S \subseteq \mathcal{A}$, and an assumption $a \in \mathcal{A}$ with $S \vdash \bar{a}$.*

- S normally attacks $\{a\}$ iff there is no $s \in S$ with $s < a$.
- $\{a\}$ reversely attacks S iff there is $s \in S$ with $s < a$.

Definition 6. *For two sets of assumptions $S, T \subseteq \mathcal{A}$,*

- S attacks T iff there are $S' \subseteq S$ and $T' \subseteq T$ such that S' normally or reversely attacks T' ;
- S is *conflict-free* iff it does not attack itself;
- S *defends* T iff S attacks each attacker of T .

We recall admissible, grounded, complete, preferred, and stable semantics for ABA^+ (abbr. *adm*, *grd*, *com*, *pref*, *stb*).

Definition 7. Let $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \neg, \leq)$ be an ABA^+ framework. Further, let $S \subseteq \mathcal{A}$ be conflict-free. Then

- $S \in adm(D)$ iff S defends itself;
- $S \in com(D)$ iff S is admissible and contains every assumption set it defends;
- $S \in grd(D)$ iff S is \subseteq -minimal in $com(D)$;
- $S \in pref(D)$ iff S is \subseteq -maximal in $com(D)$;
- $S \in stb(D)$ iff S attacks each $\{x\}$ for every $x \in \mathcal{A} \setminus S$.

For a set S of assumptions, we let $\bar{S} = \{\bar{a} \mid a \in S\}$.

Instantiating ABA Oftentimes the semantics of ABA frameworks are evaluated on a semantically equivalent abstract argumentation framework; the construction of this framework is an *instantiation* of the ABA framework (Cyras et al. 2018). We discuss a recent instantiation procedure that relies on collective attacks (Nielsen and Parsons 2006).

Definition 8. A SETAF is a pair (A, R) where A is a finite set of arguments, and $R \subseteq 2^A \times A$ is the attack relation.

SETAFs (A, R) , where for all $(T, h) \in R$ it holds that $|T| = 1$, amount to (standard Dung) AFs (Dung 1995). The semantics of SETAFs are defined in a way very similar to ABA^+ ; due to space restrictions we omit the detailed definitions. A recent overview can be found in (Flouris and Bikakis 2019; Dvořák et al. 2022).

As observed in (König, Rapberger, and Ulbricht 2022), SETAFs and ABA frameworks are closely related. We obtain a SETAF instantiation for ABA frameworks (without preferences) as follows: we construct for an ABA framework $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \neg)$ a SETAF $SF_D = (A, R)$, where the arguments A correspond to the assumptions \mathcal{A} , and for each set of assumptions S that attacks a singleton $\{a\} \subseteq \mathcal{A}$ we add an attack (S, a) to SF_D . It was shown in (König, Rapberger, and Ulbricht 2022) that now the grounded, complete, preferred, and stable extensions of D and SF_D coincide.

Example 9. Let us continue Example 4. Via the previously discussed instantiation we obtain a SETAF (left illustration) with arguments t, w , and h , and an attack from $\{w, h\}$ to t . The preference $t > w$ effectively changes the direction of the attack (right illustration), the resulting framework is no longer a SETAF and cannot be captured by the semantics proposed by Nielsen and Parsons (2006). For this reason, we generalize these semantics in the next sections.



3 Hyperframeworks

In this section, we introduce our novel formalism. We define *Hyper Argumentation Frameworks (HYPAFs)* where we allow attacks from sets of arguments to sets of arguments. We introduce HYPAFs as a general formalism that syntactically captures even more than ABA^+ . Afterwards in Section 4 we come back to HYPAFs' use as a possible instantiation for structured argumentation with preferences.

Definition 10. A HYPAF is a pair $F = (A, R)$ with a finite set of arguments A , and an attack relation $R \subseteq 2^A \times (2^A \setminus \emptyset)$.


HYPAFs (A, R) , where for all $(T, H) \in R$ it holds that $|H| = 1$ and $|T| = 1$, amount to AFs. Note that we allow for the empty set in the first position of an attack (i.e., (\emptyset, H)). The empty set in the second position of an attack however (i.e., (T, \emptyset)) we exclude. This is due to the fact that an attack towards an empty set of arguments is nonsensical and has no corresponding counter-part in any argumentation scenario.

Let us now define the concepts required to generalize the AF semantics to capture the interaction of sets of arguments.

Definition 11. Let (A, R) be a HYPAF. For $S, T \subseteq A$,

- S attacks T ($S \mapsto T$) iff there are $S' \subseteq S$ and $T' \subseteq T$ such that $(S', T') \in R$; we call S an attacker of T ;
- S is conflict-free, $S \in cf(F)$, iff it does not attack itself;
- S defends T iff for all $(U, T') \in R$ with $T' \subseteq T$, there are $S' \subseteq S$ and $U' \subseteq U$ such that $(S', U') \in R$.

Example 12. Consider the following HYPAF F with arguments a, b, c, d, e, f, g and attacks $(\{a, b\}, \{c, d, e\})$, $(\{a, b\}, \{c, d\})$, $(\{c\}, \{f\})$, and $(\{c, d\}, \{g\})$.



Both $S_1 = \{a, b, c\}$ and $S_2 = \{a, b, d\}$ are conflict-free since $\{a, b\}$ only attacks $\{c, d\}$, but none of them individually. Moreover, $\{a, b\}$ defends g , but it does not defend f .

Using these underlying notions, the definitions of the semantics naturally generalize to hyperframeworks.

Definition 13. Let $F = (A, R)$ be a HYPAF and $S \in cf(F)$. Then S is called

- admissible, $S \in adm(F)$, iff S defends itself;
- complete, $S \in com(F)$, iff $S \in adm(F)$ and S contains every set $T \subseteq A$ it defends;
- grounded, $S \in grd(F)$, iff S is \subseteq -minimal in $com(F)$;
- preferred, $S \in pref(F)$, iff S is \subseteq -maximal in $adm(F)$;
- stable, $S \in stb(F)$, iff S attacks each $T \subseteq A \setminus S$.

We observe that if a set T is defended by some set S , then all individual arguments of T are defended as well.

Lemma 14. Let $F = (A, R)$ be a HYPAF and let $S, T \subseteq A$. If S defends T then S defends $\{a\}$ for each $a \in T$.

Analogously, in order for a set S to be stable, each individual argument outside S needs to be attacked. We can thus simplify our definitions for *com* and *stb* semantics.

Lemma 15. Let $F = (A, R)$ be a HYPAF. Then

- $S \in stb(F)$ iff $S \in cf(F)$ and $S \mapsto \{x\}$ for all $x \in A \setminus S$;
- $S \in com(F)$ iff $S \in adm(F)$ and S contains each argument it defends.

Due to the first item, we note the following: If all attacks (T, H) towards some $a \in A$ satisfy $|H| \geq 2$, then $\{a\}$ cannot be defeated and hence occurs in each stable extension.

4 ABA⁺ as an Instance of HYPAFs

In this section, we state our first main result: ABA⁺ can be seen as an instance of our HYPAFs. With this, we successfully develop the first abstract instantiation for ABA⁺.

When allowing for preferences between assumptions, some attacks might be reversed. We capture these attack reversals with HYPAFs intuitively as follows.

Definition 16. *The associated HYPAF $F_D = (A, R)$ of an ABA⁺ framework $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \neg, \leq)$ is given by $A = \mathcal{A}$ and $(S, T) \in R$ iff S normally or reversely attacks T in D .*

We are ready to state our first main result: HYPAFs provide an abstract graph-based representation of ABA⁺. By design, we capture ABA⁺ semantics with our translation.


Theorem 17. *Let $\sigma \in \{adm, com, grd, pref, stb\}$ and let D be any ABA⁺ framework. It holds that $\sigma(D) = \sigma(F_D)$.*

Note that each HYPAF F_D associated to an ABA⁺ framework D satisfies that either the head or the tail has size 1; hence our translation does not capture arbitrary HYPAFs.

Observation 18. *Let D be an ABA⁺ framework, and let $F_D = (A, R)$ be the associated HYPAF. For all $(H, T) \in R$, we have either $|H| = 1$ or $|T| = 1$.*

Let us now revisit our motivating example.

Example 19. *As promised in the introduction, we can now represent Alice’s travel problem as the following HYPAF:*



We can utilize the abstract representation to compute the admissible extensions. All singletons $\{t\}$, $\{w\}$ and $\{h\}$ are unattacked, hence they are admissible; for the same reason, the sets $\{h, t\}$ and $\{w, t\}$ are admissible as well. Our theorem proves that these extensions coincide with the admissible extensions of the corresponding ABA⁺ framework.

Observation 20. *The translation from ABA⁺ to HYPAF is not bijective since different preference orderings can cause the same attack reversal. In Example 19, replacing the preference ordering $t > w$ with $t > h$ yields the same HYPAF.*

The semantics correspondence of HYPAFs and ABA⁺ established in Theorem 17 indicates certain shortcomings of our new abstract formalism. As for ABA⁺ frameworks, the existence of complete extensions is not guaranteed.

Example 21. *Let us examine complete semantics in our HYPAF from above. We start with an admissible set. As outlined before, we can accept, for instance, the set $E = \{h, t\}$ under admissible semantics. Note that w is not attacked, so w is defended by E . However, the set $A = \{h, t, w\}$ is not conflict-free. It can be checked the HYPAF has no complete extension (as expected, since it coincides with the ABA⁺ framework).*

5 Refining the HYPAF Semantics

As we have seen, the previous example illustrates certain issues. The HYPAF (and the underlying ABA⁺ framework) in fact has no complete and no grounded extension. While the absence of stable extensions is generally accepted also for AFs, the absence of complete extensions is considered

as a major drawback. After all, the 3-valued complete-based semantics serve as a compromise for reasoning in situations where no 2-valued (stable) model exists.

Since we have established an instantiation of ABA⁺ frameworks as HYPAFs, we can solve these issues by borrowing from the rich toolbox of abstract argumentation research: In this section, we recall several basic properties of AFs and study to which extend our HYPAF semantics satisfy them in our generalized setting. We will identify culprits for the absence of desirable properties and propose suitable solutions to fix these issues. Then, we will study how ABA⁺ can benefit our insights regarding HYPAFs.

5.1 General Desiderata

While stable extensions do not necessarily exist, we expect each HYPAF to possess admissible, complete, grounded, and preferred extensions. Moreover, the grounded extension should be unique since it intuitively formalizes the set of arguments that also cautious agents are willing to accept. The most important technical tool in order to ensure properties of this kind is the fundamental lemma (Dung 1995).

Lemma 22 (Fundamental Lemma, (Dung 1995)). *Let $F = (A, R)$ be an AF, $S \in adm(F)$, and $T, T' \subseteq A$ be sets of arguments that are defended by S . Then*

1. $S' = S \cup T$ is admissible, and
2. T' is defended by S' .

We collect the most important properties which are typically considered desirable for generalizing Dung’s setting.

1. (Some version of) the fundamental lemma holds.
2. There is always at least one admissible, complete, grounded, and preferred extension.
3. Every preferred extension is complete, and every stable extension is preferred.

As we will see, most of these properties are not satisfied by HYPAFs. However, let us start with the positive news: there is always at least one admissible set which also implies the existence of preferred extensions. Moreover, stable extensions are guaranteed to be preferred, as is the case for AFs.


Observation 23. *Let F be a HYPAF. Then*

1. $adm(F) \neq \emptyset$ and $pref(F) \neq \emptyset$;
2. $stb(F) \subseteq pref(F)$.

Both properties follow directly from the definitions: The empty set trivially defends itself. Moreover, stable extensions defend themselves hence they are admissible. It is also clear that they are maximal in $adm(F)$ and thus preferred.

However, the following simple example already illustrates that the semantics we defined so far violate all of the remaining properties we mentioned.

Example 24. *Let us consider the following HYPAF F :*



In F , both sets $S = \{a, b, c\}$ and $S' = \{a, b, d\}$ are conflict-free. Moreover, they are also admissible since they defend themselves and both are preferred as $\{a, b, c, d\} \notin cf(F)$.

On the other hand, what are the complete extensions of the HYPAF? We would expect that the two preferred sets $\{a, b, c\}$ and $\{a, b, d\}$ are complete as well. However, F has no complete extension at all. The reason is that S defends d since no attack is directed towards $\{d\}$ (or a subset of it). Yet, as discussed, $\{a, b, c, d\}$ is not conflict-free illustrating that the fundamental lemma is violated.

It is no surprise that the absence of the fundamental lemma causes further issues in our setting as it is essential for many AF properties. We summarize our observations.

Observation 25. Let F be a HYPAF. Then

1. the fundamental lemma is in general violated;
2. complete and grounded extensions do not always exist;
3. not every preferred extension is complete.

5.2 HYPAF Properties

The goal of this subsection is to fix these problems. To this end we will first delve into the technicalities of defense in HYPAFs. After identifying the culprits, we will propose an alternative definition for complete extensions, which ensures more well-behaved semantics.

We define the *characteristic function* for HYPAFs as it is defined for (SET)AFs: Γ_F applied to a set S of arguments returns all arguments defended by S . Due to Lemma 14 this also captures our intuition of defending sets of arguments.

Definition 26. Let $F = (A, R)$ be a HYPAF and let $S \subseteq A$. We define the characteristic function as

$$\Gamma_F(S) = \{a \in A \mid S \text{ defends } \{a\}\}.$$

Example 27. We revisit F from Example 24. Here $\Gamma_F(S) = \{a, b, c, d\}$ for each $S \subseteq A$ as each singleton is unattacked.

We mention that our characteristic function is monotonic.

Lemma 28. Let $F = (A, R)$ be a HYPAF and let $S \subseteq T \subseteq A$. Then $\Gamma_F(S) \subseteq \Gamma_F(T)$.

With the help of the characteristic function we can analyze the properties of complete-based semantics.

Complete Semantics As for AFs and SETAFs, complete semantics can be alternatively defined via the characteristic function. By definition, the complete extensions are the admissible fixed points of Γ_F .


Lemma 29. Let $F = (A, R)$ be a HYPAF. Then $S \in \text{com}(F)$ iff $S \in \text{adm}(F)$ and $S = \Gamma_F(S)$.

However, in contrast to Dung AFs and SETAFs, the characteristic function might have no admissible fixed points, as Example 27 demonstrates: $\{a, b, c, d\}$ is the only fixed point of Γ_F . We can attribute the non-existence of complete extensions to set-attacks: If the head of each attack contains at least two arguments, then complete extensions do not exist.

Lemma 30. Let $F = (A, R)$ be a HYPAF, $R \neq \emptyset$. If $|H| > 1$ for all $(T, H) \in R$ then $\text{com}(F) = \emptyset$.

Even in the special (somewhat well-behaving) case where Γ_F admits (admissible) fixed points (i.e., there are complete extensions), we are not guaranteed to have the usual semantics relations that we know from Dung’s notions.

Example 31. The following example illustrates that even if $\text{com}(F) \neq \emptyset$ holds, it is still not ensured that we also have $\text{pref}(F) \subseteq \text{com}(F)$. Consider the following HYPAF F .




We have $\emptyset \in \text{com}(F)$ and thus, $\text{com}(F) \neq \emptyset$. The set $S = \{b, c\}$ is preferred, but not complete (S defends a , but $\{a, b, c\}$ is not conflict-free). Thus $\text{pref}(F) \not\subseteq \text{com}(F)$.

We will resolve all of these problems in Section 5.3 by introducing a better suited notion of completeness.

Grounded Semantics In AFs and SETAFs, the grounded extension is the least fixed point of the characteristic function, and can be computed by via iteration, starting from the empty set until a fixed point is attained, i.e., we have $\text{grd}(F) = \{\Gamma_F^\infty(\emptyset)\}$ whenever F is a (SET)AF. For HYPAFs, this fundamental property of grounded semantics is no longer satisfied. The least fixed point of the characteristic function is not necessarily admissible or conflict-free, which causes potential non-uniqueness of the grounded extension.

Example 32. The following HYPAF F (that instantiates an ABA^+ framework) has $\text{grd}(F) = \{\{a, d, e\}, \{b, d, e\}\}$.



This result implies that skeptical reasoning w.r.t. complete semantics in HYPAFs *cannot* be done by computing the grounded extension as in (SET)AFs. However, while on the one hand we cannot guarantee that $\Gamma_F^\infty(\emptyset)$ is admissible, if this happens to be the case we are guaranteed to obtain the unique grounded extension this way.

Proposition 33. Let $F = (A, R)$ be a HYPAF. If $\Gamma_F^\infty(\emptyset) \in \text{adm}(F)$ then $\text{grd}(F) = \{\Gamma_F^\infty(\emptyset)\}$.

5.3 Revisiting the Characteristic Function

Inspired by our analysis within the previous subsection, our goal is now to refine the characteristic function Γ_F . Thereby, we do not question that Γ_F is suitable in the context of admissibility, i.e., we do not want to alter the fact that

$$\text{If } S \in \text{adm}(F), \text{ then } S \subseteq \Gamma_F(S).$$


However, it is apparent that we need a more restrictive version in order to assess whether or not some admissible set is complete. To this end recall Example 24. Here our notion of defense requires any complete extension to include all four arguments, which immediately causes a conflict. We therefore develop an alternative function Θ_F which serves to verify whether or not a set includes sufficient arguments in order to be considered “complete”. That is, we aim at:

$$\text{If } S \text{ is complete, then } \Theta_F(S) \subseteq S.$$

We refer to Θ_F as the *closure* of a set S . Our revised complete extensions need to defend themselves and be closed

under Θ_F . In order to propose such Θ_F in a natural way, let us consider the interaction between some set S and a single argument $a \in A$. In our setting, when should a belong to the “closure” of S ? A quite apparent condition is that S defends a , i.e., counters any attack directed towards a : if $(T, H) \in R$ with $a \in H$, then S attacks T . In addition, we propose a second possibility: suppose S attacks $H \setminus \{a\}$. Then (assuming we accept S), H can be considered “defeated”, but a is not involved in this defeat. Hence in this case S renders the attack (T, H) redundant and thus frees a to be acceptable. To illustrate this idea, we consider the following example.

Example 34. Let F be given as follows.



Let us argue why $\{d\}$ belongs to the closure of $S = \{g\}$. The attack $(\{e, f\}, d)$ is countered by the attack $(S, \{e, f\})$. Consider now the attack $(\{a, b\}, \{c, d\})$. Recall our interpretation of set-attacks: This attack encodes that, unless countered, $\{a, b\}$ prevents us from accepting both c and d simultaneously. However, S already defeats c , so this constraint is already satisfied given acceptance of S .

This gives the following notion of our closure operator.

Definition 35. Let $F = (A, R)$ be a HYPAF and let $S \subseteq A$. We define the closure function as

$$\Theta_F(S) = \{a \in A \mid \forall (T, H) \in R \text{ s.t. } a \in H \\ \text{it holds } S \mapsto T \text{ or } S \mapsto H \setminus \{a\}\}$$

Note that Θ_F in contrast to Γ_F considers not only attackers on $\{a\}$ but all attacks (T, H) with $a \in H$ and is in that sense thus more restrictive. However, for the additional attacks (T, H) it can also use attacks on $H \setminus \{a\}$ to protect the argument a . The more cautious Θ_F operator gives rise to an adjusted notion of complete-based semantics. Intuitively, we use our previously established notion of defense to determine admissibility, and Θ_F to assess whether further arguments have to be accepted (i.e., whether the set is “closed”).

Definition 36. Let $F = (A, R)$ be a HYPAF and let $S \subseteq A$.

- S is Θ -complete, $S \in \Theta\text{-com}(F)$, iff S is admissible and $\Theta_F(S) \subseteq S$;
- S is Θ -grounded, $S \in \Theta\text{-grd}(F)$, iff S is a \subseteq -minimal Θ -complete extension.

Example 37. Recall our F from Example 24. Let us verify that the set $S_1 = \{a, b\}$ is now Θ -complete. To this end we consider c with incoming attack $(\{a, b\}, \{c, d\})$.

- The first condition $S_1 \mapsto \{a, b\}$ does not hold;
- the second condition $S_1 \mapsto \{c, d\} \setminus \{c\}$ is also not true.

We conclude $c \notin \Theta_F(S_1)$ and analogously $d \notin \Theta_F(S_1)$. Thus $S_1 \in \Theta\text{-com}(F)$. Note that $S_2 = \{a, b, c\}$ is also Θ -complete, because $S_2 \in \text{adm}(F)$ and $d \notin \Theta_F(S_2)$. Finally, by symmetry, $S_3 = \{a, b, d\} \in \Theta\text{-com}(F)$.

In terms of faithfully generalizing the characteristic function of (SET)AFs, both Γ_F and Θ_F are equally suitable: If F is a SETAF, i.e. $|H| = 1$ for each $(T, H) \in R$, then the two notions coincide, both corresponding to the “standard” characteristic function formalizing defense in SETAFs.

Lemma 38. Let $F = (A, R)$ be a HYPAF with $|H| = 1$ for all $(T, H) \in R$. If $S \subseteq A$, then $\Gamma_F(S) = \Theta_F(S)$. In particular, $\text{com}(F) = \Theta\text{-com}(F)$ and $\text{grd}(F) = \Theta\text{-grd}(F)$.

As is the case for Γ_F , the closure function is monotonic.

Lemma 39. Let $F = (A, R)$ be a HYPAF and let $S \subseteq T \subseteq A$. Then $\Theta_F(S) \subseteq \Theta_F(T)$.

We observe that every complete extension is also Θ -complete; moreover, the Θ -grounded extension does not introduce new arguments.

Proposition 40. Let $F = (A, R)$ be a HYPAF. Then

1. $\text{com}(F) \subseteq \Theta\text{-com}(F)$;
2. if $E \in \text{grd}(F)$ exists, then $E' \in \Theta\text{-grd}(F)$ implies $E' \subseteq E$.

Recall that our motivation for introducing Θ_F was the potential non-existence of complete extensions for HYPAFs which in turn is due to the fact that Γ_F violates the fundamental lemma. For Θ_F , we can indeed derive the following result, similar in spirit to this crucial property in AFs.

Proposition 41. Let $F = (A, R)$ be a HYPAF. If $S \in \text{adm}(F)$, then $S \cup \Theta_F(S) \in \text{adm}(F)$.

This ensures that $G \in \Theta\text{-grd}(F)$ is unique and can be computed as usual: Starting from the empty set, we iterate Θ_F until reaching a fixed point. In particular this implies that there is always a complete extension (namely $\Theta_F^\infty(\emptyset)$).

Proposition 42. Let $F = (A, R)$ be a HYPAF. Then $\Theta_F^\infty(\emptyset)$ is the unique Θ -grounded extension of F . In particular, $\Theta\text{-com}(F) \neq \emptyset$ and $\Theta\text{-grd}(F) \neq \emptyset$.

Moreover, maximal $\Theta\text{-com}$ and pref extensions coincide.

Proposition 43. Let $F = (A, R)$ be a HYPAF. Then S is maximal in $\Theta\text{-com}(F)$ iff $S \in \text{pref}(F)$.

We summarize the properties of $\Theta\text{-com}$ and $\Theta\text{-grd}$.

Observation 44. Let F be a HYPAF. Then

1. Using Θ_F , a version of the fundamental lemma can be obtained;
2. $\Theta\text{-com}$ and $\Theta\text{-grd}$ extensions always exist;
3. every preferred extension is in $\Theta\text{-com}(F)$.

6 Consequences for ABA⁺

Having developed well-behaving alternatives for complete and grounded semantics, we are now in a position to reap the benefits of these results and state our Θ -semantics in the realm of ABA⁺. We close with a discussion of the complexity of ABA⁺ semantics.

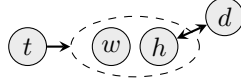
6.1 Fixing Complete-based ABA⁺ Semantics

To formalize our novel semantics, we will first define a stronger form of defense that resembles the Θ -operator. For this, we consider not only direct attacks on an assumption $a \in \mathcal{A}$, but also reversed attacks on sets that contain a .

Definition 45. Let $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \neg, \leq)$ be an ABA⁺ framework. A set $S \subseteq \mathcal{A}$ of assumptions Θ -defends $a \in \mathcal{A}$ iff for all $H, T \subseteq \mathcal{A}$ with $a \in H$, such that T normally or reversely attacks H , either (i) S attacks T ; or (ii) S attacks $H \setminus \{a\}$.

To illustrate this concept, we extend our running example.

Example 46. After carefully reviewing her budget, Alice now also considers a stay in a dormitory instead of a 5-star hotel, formalised by the ABA^+ framework and its HYPAF:

$$\begin{aligned} \mathcal{L} &= \{t, \bar{t}, d, \bar{d}, w, \bar{w}, h, \bar{h}\} \\ \mathcal{R} &= \{\bar{t} \leftarrow w, h; \bar{d} \leftarrow h; \bar{h} \leftarrow d\} \\ \mathcal{A} &= \{t, d, w, h\}, t > w, t > d, t > h \end{aligned}$$


The assumption d Θ -defends w against the attack from $\{t\}$ to $\{w, h\}$: d attacks h , hence, condition (ii) in Definition 45 applies. Note d also Θ -defends t since t is unattacked.

We are ready to give our novel definitions.

Definition 47. Let $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\cdot}, \leq)$ be an ABA^+ framework. Further, let $S \subseteq \mathcal{A}$ be conflict-free. Then

- $S \in \Theta\text{-com}(D)$ iff S is admissible and contains every assumption set it Θ -defends;
- $S \in \Theta\text{-grd}(D)$ iff S is \subseteq -minimal in $\Theta\text{-com}(D)$;

Example 48. First, let us return to our introductory example (cf. Example 1). As stated before, no complete extension exists. We apply our novel Θ -complete semantics and obtain the following extensions, capturing the intuitive outcome: $\{t\}$ (i.e., ‘Alice goes to Tokyo’), $\{t, w\}$ (‘...and stays for two weeks’) and $\{t, h\}$ (‘...or stays in a 5-star hotel’).

Coming to our example above, we note that both $\{t, w\}$ and $\{t, w, d\}$ are complete. Hence, Alice can visit Tokyo for two weeks, but cannot stay at the expensive hotel.

We can accept both sets also under Θ -complete semantics; however, now we can furthermore realize the choice between h and d . We show that we can accept $E = \{t, h\}$: indeed, E defends itself against d ; moreover, t is the only element E Θ -defends. Hence E is admissible and contains each assumption set it Θ -defends. Our Θ -complete assumption sets are $\{t\}$, $\{t, w\}$, $\{t, w, d\}$, and $\{t, h\}$. That is, with our novel semantics we also capture the possibility to have a shorter stay in Tokyo but in a 5-star hotel.

Our definitions correspond to their HYPAF counter-parts.

Theorem 49. Let $\sigma \in \{\Theta\text{-com}, \Theta\text{-grd}\}$ and let D be any ABA^+ framework. Then $\sigma(D) = \sigma(F_D)$.

Together with the results from the previous section we can now derive several crucial properties of our Θ -semantics:

1. With our notion of Θ -defense (cf. Definition 45), we can entail a version of the fundamental lemma;
2. $\Theta\text{-com}$ and $\Theta\text{-grd}$ extensions always exist;
3. every preferred extension is Θ -complete.

Moreover, each complete extension is Θ -complete; also, the Θ -grounded extension is unique and contained in the grounded extension (whenever the latter exists). With our novel semantics we therefore provide faithful generalizations of both semantics for ABA^+ that are guaranteed to return an extension while preserving their original spirit.

6.2 Computational Complexity

Finally, we consider the computational complexity of our novel semantics and compare them with the complexity of the existing semantics. We also provide a refined analysis

	<i>adm</i>	<i>grd</i>	Θ - <i>grd</i>	<i>com</i>	Θ - <i>com</i>	<i>pref</i>	<i>stb</i>
Ver_σ	coNP-c	Π_2^P -c	DP-c	DP-c	DP-c	in Π_2^P	in P
$Cred_\sigma$	Σ_2^P -c	in Σ_3^P	in Δ_2^P	Σ_2^P -c	Σ_2^P -c	Σ_2^P -c	NP-c
$Skept_\sigma$	triv.	Π_2^P -c	in Δ_2^P	Π_2^P -c	in Δ_2^P	in Π_3^P	coNP-c

Table 1: Complexity of ABA^+ semantics.

of the traditional ABA^+ semantics, extending the results from (Lehtonen, Wallner, and Järvisalo 2021a,b). We assume familiarity with the polynomial hierarchy. In our analysis, we focus on the classical reasoning tasks: *verification* of an extension set under semantics σ (Ver_σ) as well as *credulous* ($Cred_\sigma$) and *skeptical acceptance* ($Skept_\sigma$), i.e., deciding whether a given conclusion is contained in some or each σ -extension, resp. (cf. (Dvořák and Dunne 2017)).

Theorem 50. The complexity results in Table 1 hold.

The results for *adm* and *stb* are due to Lehtonen, Wallner, and Järvisalo (2021a). We extend the results of (Lehtonen, Wallner, and Järvisalo 2021a,b) for *grd*, *com*, and *pref*.

Key to establishing the results for our Θ -semantics is the complexity of Θ -defense. We show that the problem is NP-complete, which is in accordance with the complexity of deciding standard defense (Lehtonen, Wallner, and Järvisalo 2021a). With this at hand we can show that verification and credulous acceptance for *com* and Θ -*com* coincides. Since the Θ -grounded extension can be computed via fixed point iteration, the complexity of Θ -*grd* semantics and $Skept_{\Theta\text{-com}}$ is less demanding, compared to their traditional variants.

7 Discussion and Related Work

We tackled two crucial drawbacks of ABA^+ : the missing abstract instantiation and the problem of non-existence of grounded and complete extensions. To this end, we introduced Hyper Argumentation Frameworks, based on set-to-set attacks, and showed that ABA^+ is an instance of the novel formalism. We developed faithful generalizations of complete and grounded semantics which are guaranteed to return an output, without causing a rise in complexity.

While set-to-set attacks as used in HYPAFs received only limited attention so far, the idea of *collective* attack appeared in different variants in the KR community, the most prominent example being SETAFs (Nielsen and Parsons 2006). They discuss potential ways to capture set-to-set attacks via SETAFs. Verheij (1996) models defeat not via directed graphs but using rule-like statements; Bochman (2003) uses similar concepts to formalize global conflicts; Nielsen and Parsons (2006) reduce set-to-set attacks to SETAFs; and Gabbay and Gabbay (2016) investigate (among other notions) cases where the attacking set applies conjunctively and the attacked set is understood disjunctively. An overview of possible interpretations of set-to-set attacks is given in the workshop paper (Dimopoulos et al. 2023), where attacks are also interpreted in an indeterministic setting.

Future work includes further investigating our newly established Θ -semantics in the context of applications of ABA^+ and other structured formalisms as well as settling the remaining tight bounds in our complexity analysis.

Acknowledgments

This work was funded by the Austrian Science Fund (FWF) under grant P32830, by the Vienna Science and Technology Fund (WWTF) under grant ICT19-065, 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, and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101020934).

References

- Amgoud, L.; and Vesic, S. 2014. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.*, 55(2): 585–606.
- Bao, Z.; Cyras, K.; and Toni, F. 2017. ABAPlus: Attack Reversal in Abstract and Structured Argumentation with Preferences. In An, B.; Bazzan, A. L. C.; Leite, J.; Villata, S.; and van der Torre, L. W. N., eds., *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference, Nice, France, October 30 - November 3, 2017, Proceedings*, 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.
- Bench-Capon, T. J. M.; Prakken, H.; and Sartor, G. 2009. Argumentation in Legal Reasoning. In Simari, G. R.; and Rahwan, I., eds., *Argumentation in Artificial Intelligence*, 363–382. Springer.
- Besnard, P.; and Hunter, A. 2014. Constructing argument graphs with deductive arguments: a tutorial. *Argument Comput.*, 5(1): 5–30.
- Bochman, A. 2003. Collective Argumentation and Disjunctive Logic Programming. *J. Log. Comput.*, 13(3): 405–428.
- Caminada, M.; Sá, S.; Alcântara, J.; and Dvořák, W. 2015. On the Difference between Assumption-Based Argumentation and Abstract Argumentation. *IfCoLog Journal of Logic and its Applications*, 2(1): 15–34.
- Cyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2018. Assumption-Based Argumentation: Disputes, Explanations, Preferences. In Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds., *Handbook of Formal Argumentation*, chapter 7, 365–408. College Publications. Also appears in *IfCoLog Journal of Logics and their Applications* 4(8):2407–2456.
- Cyras, K.; and Toni, F. 2016. ABA+: Assumption-Based Argumentation with Preferences. In Baral, C.; Delgrande, J. P.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016*, 553–556. AAAI Press.
- Dimopoulos, Y.; Dvořák, W.; König, M.; Rapberger, A.; Ulbricht, M.; and Woltran, S. 2023. Sets Attacking Sets in Abstract Argumentation. In Sauerwald, K.; and Thimm, M., eds., *Proceedings of the 21st International Workshop on Non-Monotonic Reasoning co-located with the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR 2023) and co-located with the 36th International Workshop on Description Logics (DL 2023), Rhodes, Greece, September 2-4, 2023*, volume 3464 of *CEUR Workshop Proceedings*, 22–31. CEUR-WS.org.
- 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): 321–358.
- Dvořák, W.; and Dunne, P. E. 2017. Computational Problems in Formal Argumentation and their Complexity. *FLAP*, 4(8): 2557–2622.
- Dvořák, W.; König, M.; Ulbricht, M.; and Woltran, S. 2022. Rediscovering Argumentation Principles Utilizing Collective Attacks. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning KR22*, 122–131.
- Flouris, G.; and Bikakis, A. 2019. A comprehensive study of argumentation frameworks with sets of attacking arguments. *Int. J. Approx. Reason.*, 109: 55–86.
- Gabbay, D. M.; and Gabbay, M. 2016. Theory of disjunctive attacks, Part I. *Log. J. IGPL*, 24(2): 186–218.
- König, M.; Rapberger, A.; and Ulbricht, M. 2022. Just a Matter of Perspective. In Toni, F.; Polberg, S.; Booth, R.; Caminada, M.; and Kido, H., eds., *Computational Models of Argument - Proceedings of COMMA 2022, Cardiff, Wales, UK, 14-16 September 2022*, volume 353 of *Frontiers in Artificial Intelligence and Applications*, 212–223. IOS Press.
- Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2021a. Declarative Algorithms and Complexity Results for Assumption-Based Argumentation. *J. Artif. Intell. Res.*, 71: 265–318.
- Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2021b. Harnessing Incremental Answer Set Solving for Reasoning in Assumption-Based Argumentation. *Theory Pract. Log. Program.*, 21(6): 717–734.
- Modgil, S.; and Prakken, H. 2014. The ASPIC⁺ framework for structured argumentation: a tutorial. *Argument & Computation*, 5(1): 31–62.
- 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 2006*, 54–73. Springer.
- Verheij, H. 1996. *Rules, reasons, arguments : formal studies of argumentation and defeat*. Ph.D. thesis, Maastricht University, Netherlands.