On the Expressive Power of Assumption-Based Argumentation

Matti Berthold[1(✉)], Anna Rapberger[2(✉)], and Markus Ulbricht[1(✉)]

[1] ScaDS.AI Dresden/Leipzig, Universität Leipzig, Leipzig, Germany
{berthold,ulbricht}@informatik.uni-leipzig.de
[2] TU Wien, Vienna, Austria
anna.rapberger@tuwien.ac.at

Abstract. The expressiveness of any given formalism lays the theoretical foundation for more specialized topics such as investigating dynamic reasoning environments. The modeling capabilities of the formalism under investigation yield immediate (im)possibility results in such contexts. In this paper we investigate the expressiveness of assumption-based argumentation (ABA), one of the major structured argumentation formalisms. In particular, we examine so-called *signatures*, i.e., sets of extensions that can be realized under a given semantics. We characterize the signatures of common ABA semantics for flat, finite frameworks with and without preferences. We also give several results regarding conclusion-based semantics for ABA.

1 Introduction

Within the last decades, AI research has witnessed an increasing demand for knowledge representation systems that are capable of handling inconsistent beliefs. Research in computational argumentation has addressed this issue by developing numerous sophisticated methods to representing and analyzing conflicting information [22]. A key player in this field are abstract argumentation frameworks (AFs) as proposed by Dung in 1995 [15]. In AFs, arguments are interpreted as atomic entities and conflicts as a binary relation; consequently, an AF represents a given debate as a directed graph F. Research on AFs is driven by various *semantics* which strive to formalize what reasonable viewpoints F entails. That is, if $E \in \sigma(F)$ for a semantics σ, then E is interpreted as a jointly acceptable set of arguments. These sets E are called *σ-extensions* of F.

In the research area of *structured argumentation*, an AF is constructed from a given knowledge base in order to explicate arising conflicts in a comprehensible graph. One highly influential approach in this area is *assumption-based argumentation* (ABA) [8,12]. *Assumptions* provide the foundation for arguments and determine their conflicts. ABA frameworks (ABAFs) are also evaluated under so-called semantics; in contrast to many other argumentation formalisms, the native ABA semantics output sets of assumptions rather than arguments.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
S. Gaggl et al. (Eds.): JELIA 2023, LNAI 14281, pp. 145–160, 2023.
https://doi.org/10.1007/978-3-031-43619-2_11

Within the last years, researchers have studied the modeling capabilities of different AF semantics extensively [6,16,33]. To this end the notion of the *signature* Σ_σ of a semantics σ has been coined. This concept formalizes what can be the set of σ-extensions of an AF, i.e., $\Sigma_\sigma = \{\sigma(F) \mid F \text{ is a finite AF}\}$. Some properties of important semantics are folklore within the AF community. For example, the empty set is always *admissible* and the *stable* extensions of an AF are incomparable[1]. However, establishing the precise characterizations for the common AF semantics is a challenging endeavor [16].

The signatures of argumentation semantics are an important formal tool underlying several applications as well as theoretical results building upon them. Recent years witnessed significant developments in the construction of explanations based on formal argumentation [13,34]. Key to obtain argumentative explanations are translations of the given (rule-based) knowledge base into a suitable abstract argumentation formalism [21]. Such formalisms differ in their expressive power and thus in their ability to provide semantics-preserving translations. Signature characterizations for different abstract and structured formalisms thus pave the way for developing suitable translations, facilitating the extraction of argumentative explanations. Precise characterizations of the modeling capacities of semantics furthermore play a central role in the context of dynamic reasoning environments, i.e., knowledge bases that evolve over time [22]. Many research questions on dynamics heavily rely on insights as to how the models of a given AF can be manipulated in order to reach a certain goal. A noteworthy example is the current hot topic of forgetting [2,5,7,25] where the goal is oftentimes to cut arguments out of or remove extensions entirely. Whether the target modification is attainable can be decided by studying the signatures of the semantics.

While signatures have been investigated for various abstract argumentation formalisms [16,17,20], this line of research has mostly been neglected in the realm of structured argumentation. In this paper, we tackle this issue and present various results regarding the expressive power of ABA. We first consider the most common ABA fragment and fully characterize the signatures of all standard semantics commonly studied in the literature. We achieve this by building upon previous results from abstract argumentation research. We then study various aspects, adding to our investigation by shifting the focus to the conclusions of the extensions or incorporating preferences.

2 Background

We recall assumption-based argumentation (ABA) [12], argumentation frameworks with collective attacks (SETAFs) [27], and their relation [24].

Assumption-Based Argumentation. We consider 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.

[1] We refer to Sect. 2 for a formal introduction of the semantics we consider.

Definition 1. *An ABA framework (ABAF) is a tuple* $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$*, where* $(\mathcal{L}, \mathcal{R})$ *is a deductive system,* $\mathcal{A} \subseteq \mathcal{L}$ *a set of assumptions, and* $\overline{} : \mathcal{A} \to \mathcal{L}$ *a contrary function.*

In this work, we focus on frameworks which are *flat*, i.e., $head(r) \notin \mathcal{A}$ for each rule $r \in \mathcal{R}$, and *finite*, i.e., $\mathcal{L}$, $\mathcal{R}$, $\mathcal{A}$ are finite. By $\mathcal{A}(\mathcal{D})$ and $\mathcal{L}(\mathcal{D})$ we denote the assumptions and atoms occurring in $\mathcal{D}$, respectively.

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 $S \vdash_R p$. Moreover, we let $Th_{\mathcal{D}}(S) = \{p \in \mathcal{L} \mid S \vdash p\}$.

A set of assumptions S *attacks* a set of assumptions T if there are $S' \subseteq S$ and $a \in T$ *s.t.* $S' \vdash \overline{a}$; The set S is conflict-free ($S \in cf(\mathcal{D})$) if it does not attack itself; S *defends* $a \in \mathcal{A}$ if for each attacker T of $\{a\}$, we have S attacks T. A conflict-free set S is admissible ($S \in ad(\mathcal{D})$) iff S defends each $a \in S$. We recall grounded, complete, preferred, and stable ABA semantics (abbr. gr, co, pr, stb).

Definition 2. *Let* $\mathcal{D}$ *be an ABAF and let* $S \in ad(\mathcal{D})$*. Then*

- $S \in co(\mathcal{D})$ *iff* S *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 *attacks each* $\{x\} \subseteq \mathcal{A}(\mathcal{D}) \setminus S$*.*

Example 1. We consider an ABAF $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ with $\mathcal{L} = \{a, b, c, a_c, b_c, c_c\}$, assumptions $\mathcal{A} = \{a, b, c\}$, their contraries a_c, b_c, and c_c, respectively, and rules

$$a_c \leftarrow b, c \qquad\qquad b_c \leftarrow a \qquad\qquad c_c \leftarrow a, b$$

Then the set $\{a\}$ is admissible: it defends itself against its only attacker $\{b, c\}$, by attacking b. The set $\{a\}$ is not complete, however, since it also defends the assumption c. The sets $\{a, c\}$ and $\{b, c\}$ are complete, preferred and stable. Moreover, $\emptyset$ is complete and the unique grounded extension of $\mathcal{D}$. ◊

SETAFs. We recall *argumentation frameworks with collective attacks* [27].

Definition 3. *A SETAF is a pair* $F = (A, R)$ *where* A *is a finite set of arguments and* $R \subseteq (2^A \setminus \{\emptyset\}) \times A$ *encodes attacks.*

SETAFs generalize Dung's *abstract argumentation frameworks (AFs)* [15]. In AFs, each attacking set is a singleton, i.e., $|T| = 1$ for each $(T, h) \in R$. The SETAF semantics are defined in a way that they naturally generalize Dung's AF semantics. They are, however, even closer in spirit to ABA semantics.

A set of arguments S *attacks* an argument $a \in A$ if there is some $S' \subseteq S$ such that $(S', a) \in R$; S attacks a set of arguments T if there are $S' \subseteq S$ and $t \in T$ such that $(S', t) \in R$; S is conflict-free ($S \in cf(F)$) if it does not attack itself.

A set S *defends* an argument $a \in A$ if for each attacker T of a, it holds that S attacks T; S defends $T \subseteq A$ iff it defends each $t \in T$. A conflict-free set S is admissible ($S \in ad(F)$) iff S defends each $a \in S$. We recall grounded, complete, preferred, and stable SETAF semantics (abbr. gr, co, pr, and stb).

Definition 4. *Let F be a SETAF and let $S \in ad(F)$. Then,*

- *$S \in co(F)$ iff S contains each argument it defends;*
- *$S \in gr(F)$ iff S is $\subseteq$-minimal in $co(F)$;*
- *$S \in pr(F)$ iff S is $\subseteq$-maximal in $ad(F)$;*
- *$S \in stb(F)$ iff S attacks all $a \in A(F) \setminus S$.*

Relating ABAFs and SETAFs. For our first main result we exploit the close connection of ABAFs and SETAFs. The key idea is to identify assumptions in ABAFs with arguments in SETAFs; moreover, attacks between assumption-sets can be viewed as collective attacks between arguments in SETAFs and vice versa. The following translations are due to [24].

Definition 5. *For an ABAF $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$, the corresponding SETAF $F_\mathcal{D} = (A_\mathcal{D}, R_\mathcal{D})$ is defined by $A_\mathcal{D} = \mathcal{A} \setminus \{a \mid \overline{a} \in Th_\mathcal{D}(\emptyset)\}$ and for $S \cup \{a\} \subseteq A_\mathcal{D}$ we let $(S, a) \in R_\mathcal{D}$ iff $S \vdash \overline{a}$.*[2] *For a SETAF $F = (A, R)$, the corresponding ABAF $\mathcal{D}_F = (\mathcal{L}_F, \mathcal{R}_F, \mathcal{A}_F, \overline{})$ is defined by $\mathcal{L}_F = A \cup \{p_x \mid x \in A\}$, $\mathcal{A}_F = A$, $\overline{x} = p_x$ for all $x \in A$, and for each $(T, h) \in R$, we add a rule $p_h \leftarrow T$ to $\mathcal{R}_F$.*

Example 2. Consider the ABAF $\mathcal{D}$ from Example 1. The corresponding SETAF $F_\mathcal{D}$ has the arguments $A_\mathcal{D} = \{a, b, c\}$; moreover, the arguments determine the collective attacks. For instance, from $\{b, c\} \vdash a_c$ we obtain that $\{b, c\}$ collectively attacks a. Below, we depict all attacks between the assumption-sets as usually done in the literature (left) and the corresponding SETAF (right). Left, we omit the (irrelevant) $\emptyset$ and (self-attacking) $\mathcal{A}$.

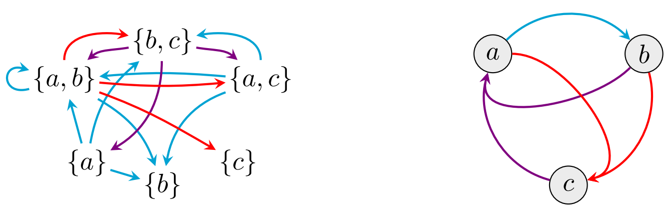

Attacks obtained from $\{a\} \vdash b_c$ are in cyan, from $\{b, c\} \vdash a_c$ in violet, and attacks obtained from $\{a, b\} \vdash c_c$ are depicted in red. Overall, we observe that the SETAF representation is significantly smaller: in contrast to the traditional ABA set representation, it requires only a single node for each assumption. ◊

We recall the close relation between ABAFs and SETAFs.

Proposition 1. *Given a semantics $\sigma \in \{ad, gr, co, pr, stb\}$. For an ABAF $\mathcal{D}$ and its associated SETAF $F_\mathcal{D}$, it holds that $\sigma(\mathcal{D}) = \sigma(F_\mathcal{D})$. For a SETAF F and its associated ABAF $\mathcal{D}_F$, it holds that $\sigma(F) = \sigma(\mathcal{D}_F)$.*

[2] We note that the original translation slightly deviates from this version.

In [24], the result has only been stated for gr, co, pr, stb semantics; however, the adaption to admissible semantics can be easily obtained.

3 Signatures of ABA Frameworks

The investigation of the *signature* of a semantics is driven by properties of a given set $\mathbb{S}$ of sets, in order to assess whether it is conceivable that there is some knowledge base (in our case: some ABAF) $\mathcal{D}$ *s.t.* $\sigma(\mathcal{D}) = \mathbb{S}$. Let us familiarize with this setting by considering the following example.

Example 3. Let $\mathbb{S} = \{\{a\}, \{a, b\}, \{a, c\}\}$. We can actually already infer a lot about this set.

- It is impossible that $\mathbb{S}$ corresponds to the ad sets of an ABA knowledge base, i.e., $ad(\mathcal{D}) = \mathbb{S}$; the reason is that $\emptyset \in ad(\mathcal{D})$ for any ABAF $\mathcal{D}$, but $\emptyset \notin \mathbb{S}$;
- $\mathbb{S}$ cannot correspond to gr since $|gr(\mathcal{D})| = 1$ for any ABAF $\mathcal{D}$;
- it is also impossible that $stb(\mathcal{D}) = \mathbb{S}$ or $pr(\mathcal{D}) = \mathbb{S}$, because stable and preferred sets are always incomparable; however, in $\mathbb{S}$ we have $\{a\} \subsetneq \{a, b\}$;
- it is however possible to construct $\mathcal{D}$ with $co(\mathcal{D}) = \mathbb{S}$. The set $\{a\}$ could be the grounded extension and b and c in a mutual attack, yielding the complete extensions $co(\mathcal{D}) = \{\{a\}, \{a, b\}, \{a, c\}\}$. ♢

We now formally define ABA signatures.

Definition 6. *Given a semantics σ, the* signature *of σ is*

$$\Sigma_\sigma^{ABA} = \{\sigma(\mathcal{D}) \mid \mathcal{D} \textit{ is a flat, finite ABAF}\}.$$

Signatures are sets of sets of assumptions, i.e., $\Sigma_\sigma^{ABA} \subseteq 2^{2^{\mathcal{U}}}$ where $\mathcal{U}$ denotes the set of all possible (countably infinitely many) assumptions. We call a set $\mathbb{S} \subseteq 2^{\mathcal{U}}$ an *extension-set*. An extension-set $\mathbb{S}$ is *realizable* under the given semantics σ, if there exists a ABAF $\mathcal{D}$ that *realizes* it, i.e., $\sigma(\mathcal{D}) = \mathbb{S}$.

We will infer Σ_σ^{ABA} by exploiting the close relation to SETAFs. To this end we recall the concept of their signatures, given as

$$\Sigma_\sigma^{SF} = \{\sigma(F) \mid F \text{ is a SETAF}\}.$$

Analogously, signatures for SETAFs are sets of sets of arguments. The concepts of *extension-sets* and *realizations* naturally transfer to this setting.

We are now ready to study the ABA signatures. Before we can delve into our results, however, we need to introduce some theoretical machinery (cf. [16,17]).

Definition 7. *Let $\mathbb{S}$ be a set of sets. We let*

$$\mathbb{A}_\mathbb{S} = \bigcup \mathbb{S}, \quad \mathbb{P}_\mathbb{S} = \{S \subseteq \bigcup \mathbb{S} \mid \nexists S' \in \mathbb{S} : S \subseteq S'\}, \quad dcl(\mathbb{T}) = \{S' \subseteq S \mid S \in \mathbb{T}\}.$$

Thereby, $\mathbb{P}_\mathbb{S}$ is the set of potential conflicts *in $\mathbb{S}$ and $dcl(\mathbb{T})$ the* downward closure *of $\mathbb{T}$. The* completion-sets *of a set of assumptions T in $\mathbb{S}$ are given by*

$$\mathbb{C}_\mathbb{S}(T) = \{S \in \mathbb{S} \mid T \subseteq S,\ \nexists S' \in \mathbb{S} : T \subseteq S' \subsetneq S\}.$$

Let us illustrate these concepts in the following example.

Example 4. Let $\mathbb{S} = \{\{a\}, \{a, b\}, \{a, c\}\}$. We have the following sets:

- $\mathbb{A}_{\mathbb{S}} = \{a, b, c\}$ intuitively corresponding to credulously accepted assumptions;
- $\mathbb{P}_{\mathbb{S}} = \{\{b, c\}, \{a, b, c\}\}$ since b and c never occur in the same extension;
- $\mathbb{C}_{\mathbb{S}}(\{a\})$ is only $\{a\}$ itself; on the other hand, $\mathbb{C}_{\mathbb{S}}(\{b\}) = \{\{a, b\}\}$ since $\{a, b\}$ is a (the only) minimal set containing $\{b\}$.
- The downward closure of $\mathbb{S}$ is the set $dcl(\mathbb{S}) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}\}$. Intuitively, if $\mathbb{S}$ are the extensions of some ABA $\mathcal{D}$, we can be certain that each set in $dcl(\mathbb{S})$ if conflict-free. ♢

Having established the sets we require, let us now consider relevant properties.

Definition 8. *Given an ABAF $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$. A set $\mathbb{S} \subseteq 2^{\mathcal{A}}$ is*

- incomparable *if for $S, S' \in \mathbb{S}$, $S \subseteq S'$ implies $S = S'$;*
- set-conflict-sensitive *if for all $S, S' \in \mathbb{S}$ with $S \cup S' \notin \mathbb{S}$ it holds that there is some $p \in S$ such that $S' \cup \{p\} \in \mathbb{P}_{\mathbb{S}}$;*
- set-com-closed *if for all $\mathbb{T}, \mathbb{U} \subseteq \mathbb{S}$, the following holds: if their elements $T = \bigcup \mathbb{T}$ and $U = \bigcup \mathbb{U}$ are both contained in the downward closure of $\mathbb{S}$ and satisfy $|\mathbb{C}_{\mathbb{S}}(T \cup U)| \neq 1$ then there is an assumption $u \in U$ such that $T \cup \{u\} \in \mathbb{P}_{\mathbb{S}}$.*

Example 5. We continue the above example with $\mathbb{S} = \{\{a\}, \{a, b\}, \{a, c\}\}$:

- $\mathbb{S}$ is not incomparable since $\{a\} \subsetneq \{a, b\}$;
- $\mathbb{S}$ is set-conflict-sensitive. The only sets with $S, S' \in \mathbb{S}$ with $S \cup S' \notin \mathbb{S}$ are $\{a, b\}$ and $\{a, c\}$. Now consider $b \in S$. Indeed, $S' \cup \{b\} = \{a, b, c\} \in \mathbb{P}_{\mathbb{S}}$. Intuitively, this formalizes that the union $S \cup S'$ is not an extension, i.e., not contained in $\mathbb{S}$, since b and c cause a conflict.
- $\mathbb{S}$ is set-com-closed. Take for example $\mathbb{T} = \{\{a\}, \{a, b\}\}$ and $\mathbb{U} = \{\{a, c\}\}$. We thus have $T = \{a, b\}$ and $U = \{a, c\}$. Both T and U are contained in the downward closure $dcl(\mathbb{S})$ we calculated before. For the union $T \cup U = \{a, b, c\}$ we have $\mathbb{C}_{\mathbb{S}}(T \cup U) = \emptyset$ since no superset of $T \cup U$ occurs in $\mathbb{S}$. Therefore, the condition $|\mathbb{C}_{\mathbb{S}}(T \cup U)| \neq 1$ fires and we need to find $u \in U$ *s.t.* $T \cup \{u\} \in \mathbb{P}_{\mathbb{S}}$. Indeed, c occurs in U and $T \cup \{c\} = \{a, b, c\}$ is a potential conflict.
  The rationale behind this property is the following: Suppose we consider complete semantics. Then, $\mathbb{C}_{\mathbb{S}}(\emptyset)$ is the grounded extension and we thus have $|\mathbb{C}_{\mathbb{S}}(\emptyset)| = 1$. This does not only apply to the empty set; given some admissible extension E, there is also always a unique minimal complete extension containing E. The set-com-closed property extracts situations where $|\mathbb{C}_{\mathbb{S}}(\cdot)| = 1$ must hold; if not, then we need to find a corresponding "reason", i.e., some $u \in U$ causing the conflict, i.e., $T \cup \{u\} \in \mathbb{P}_{\mathbb{S}}$. ♢

We are ready to state the main result of this section.

Theorem 1. *It holds that*

$$\begin{aligned}
\Sigma_{gr}^{ABA} &= \{\mathbb{S} \mid |\mathbb{S}| = 1\},\\
\Sigma_{ad}^{ABA} &= \{\mathbb{S} \neq \emptyset \mid \mathbb{S} \textit{ is set-conflict-sensitive and } \emptyset \in \mathbb{S}\},\\
\Sigma_{co}^{ABA} &= \{\mathbb{S} \neq \emptyset \mid \mathbb{S} \textit{ is set-com-closed and } \bigcap \mathbb{S} \in \mathbb{S}\},\\
\Sigma_{stb}^{ABA} &= \{\mathbb{S} \mid \mathbb{S} \textit{ is incomparable}\}, \textit{and}\\
\Sigma_{pr}^{ABA} &= \Sigma_{stb}^{ABA} \setminus \{\emptyset\}.
\end{aligned}$$

We obtain the result by (1) exploiting the connection between SETAFs and ABAFs [24] (cf. Definition 5 and Proposition 1) in order to (2) transfer signature results for SETAF semantics [17] to the associated ABAF semantics.

4 Compact Realizability in ABA

In the previous section we could establish the plain ABA signatures by exploiting the close relation to SETAFs. In the remainder of this paper we will study further aspects which require more specialized techniques. In the context of AF signatures it was observed that there are extension-sets that can only be realized by the use of auxiliary arguments that are never accepted. An AF F is *compact w.r.t. a semantics* σ iff each argument in F is credulously accepted [4]. This notion can be translated to ABA and be employed to prove certain unsatisfiability results.

Definition 9. *Given* σ*, an ABAF* $\mathcal{D}$ *is* compact *w.r.t.* σ *iff* $\mathcal{A}(\mathcal{D}) = \bigcup \sigma(\mathcal{D})$.

We term a semantics σ to be *compactly realizable*, iff for any $\mathcal{D}$ there exists a $\mathcal{D}'$ that is equivalent to $\mathcal{D}$ under σ such that $\mathcal{D}'$ is compact. In the remainder of this section, we prove the following theorem.

Theorem 2. *For ABA, the semantics gr and pr are compactly realizable, whereas ad and co are not. The semantics stb is compactly realizable if we limit ourselves to non-empty extension-sets.*

The empty extension-set $\mathbb{S} = \emptyset$ is not compactly realizable under stable semantics since an ABAF with no assumption has the unique stable extension $\emptyset$. To prove the compact realizability of *gr* and *pr*, and *stb* (for non-empty extension-sets) we employ canonical constructions for ABAFs[3] that are similar in spirit to SETAF constructions for these semantics [17]. We first show compactness of *gr* semantics. For this, we construct a canonical ABA $\mathcal{D}$ with no rule at all.

Definition 10. *Given an extension-set* $\mathbb{S}$ *with* $|\mathbb{S}| = 1$*, i.e.,* $\mathbb{S} = \{S\}$*, we let* $\mathcal{D}_{\mathbb{S}}^{gr} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$*, where* $\mathcal{L} = \mathcal{A} \cup \{a_c \mid a \in \mathcal{A}\}$*,* $\mathcal{R} = \emptyset$*,* $\mathcal{A} = S$*, and* $\overline{a} = a_c$ *for each* $a \in \mathcal{A}$*.*

[3] Implementation of the canonical constructions for all semantics considered in this paper are available at https://pyarg.npai.science.uu.nl/ [28].

It is easy to see that $gr(\mathcal{D}_{\mathbb{S}}^{gr}) = \{S\}$ if $\{S\} = \mathbb{S}$; thus, this construction realizes $\mathbb{S}$ under *gr* semantics. For *pr* and *stb* semantics, we proceed as follows.

Definition 11. *Given an incomparable, non-empty extension-set* $\mathbb{S}$, *we define* $\mathcal{D}_{\mathbb{S}}^{inc} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$, *where* $\mathcal{L} = \mathcal{A} \cup \{a_c \mid a \in \mathcal{A}\}$, $\mathcal{R} = \{a_c \leftarrow S \mid a \notin S, S \in \mathbb{S}\}$, $\mathcal{A} = \bigcup \mathbb{S}$, *and* $\overline{a} = a_c$ *for each* $a \in \mathcal{A}$.

Example 6. Let $\mathbb{S} = \{\{a, b\}, \{a, c\}, \{b, c\}\}$. We construct the ABAF $\mathcal{D}$ with assumptions $\mathcal{A} = \{a, b, c\}$ and rules $\mathcal{R} = \{a_c \leftarrow b, c;\ b_c \leftarrow a, c;\ c_c \leftarrow a, b\}$. Note that $\mathcal{L}$ and $\overline{}$ are now also determined. Indeed, $\mathcal{D}$ realizes our desired assumption-set under stable and preferred semantics, e.g., $\sigma(\mathcal{D}) = \{\{a, b\}, \{a, c\}, \{b, c\}\}$. ♢

The compact realizability of *gr*, *pr* and *stb* holds due to the constructions in Definitions 10 and 11 not employing unaccepted assumptions.

Proposition 2. *The semantics gr and pr are compact realizable. The semantics stb is compactly realizable if we limit ourselves to non-empty extension-sets.*

In contrast, admissible and complete semantics are not compact realizable. We give corresponding counter-examples.

Example 7. (admissible) Consider $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ with

$$\mathcal{L} = \{a, b, c, a_c, b_c, c_c\} \qquad \mathcal{R} = \{a_c \leftarrow b;\ b_c \leftarrow c\}$$
$$\mathcal{A} = \{a, b, c\} \qquad \overline{} = \{(a, a_c), (b, b_c), (c, c_c)\}$$

Then $ad(\mathcal{D}) = \{\emptyset, \{c\}, \{a, c\}\}$, but there is no ABAF $\mathcal{D}'$ with $\mathcal{A}(\mathcal{D}') = \{a, c\}$, *s.t.* $ad(\mathcal{D}') = \{\emptyset, \{c\}, \{a, c\}\}$. It is impossible to express c supporting a without the presence of a third assumption.

(complete). Consider $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$, with

$$\mathcal{L} = \{a, b, a_c, b_c\} \qquad \mathcal{R} = \{b_c \leftarrow a;\ a_c \leftarrow b;\ a_c \leftarrow a\}$$
$$\mathcal{A} = \{a, b\} \qquad \overline{} = \{(a, a_c), (b, b_c)\}$$

Then $co(\mathcal{D}) = \{\emptyset, \{b\}\}$, but there is no ABAF $\mathcal{D}'$ with $\mathcal{A}(\mathcal{D}') = \{b\}$, *s.t.* $co(\mathcal{D}') = \{\emptyset, \{b\}\}$, because if b is the only assumption, then there is only one complete extension as there is nothing against b could defend itself. ♢

From these observations, Theorem 2 follows.

5 Claims, Preferences and Beyond

So far, we put our focus on the most common ABA fragment. In particular, we considered semantics in terms of assumption-sets only. There are, however, several other aspects of ABA that can also be taken into account. In this section, we present several results for realizing extension sets in ABA$^+$ which extends the basic setting by allowing for preferences between assumptions. Moreover, we outline some insights regarding signatures for conclusion sets, i.e., we evaluate $\mathcal{D}$ in terms of accepted conclusions, not just the underlying assumptions.

Signatures for Conclusion Extensions

Let us now also consider the set of all atoms derivable from an assumption-set. Recall that by

$$Th_{\mathcal{D}}(S) = \{p \mid \exists S' \subseteq S : S' \vdash p\}$$

we denote the set of all conclusions derivable from an assumption-set S in an ABAF $\mathcal{D}$. Observe that $S \subseteq Th_{\mathcal{D}}(S)$ since per definition, each assumption $a \in \mathcal{A}$ is derivable from $\{a\} \vdash_{\emptyset} a$. We define the conclusion-based semantics for ABA by considering the derivable conclusions of acceptable assumption sets.

Definition 12. *Let $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ be an ABAF and let σ be a semantics. We define its* conclusion-based variant *as $\sigma_c(\mathcal{D}) = \{Th_{\mathcal{D}}(E) \mid E \in \sigma(\mathcal{D})\}$.*

We write $\Sigma_{\sigma_c}^{ABA}$ to denote the conclusion-based signatures. In this section, we compare the conclusion-based signatures with the standard ABA signatures from above. Since deriving the conclusions as well gives more fine-grained extensions, the attentive reader might expect that this setting is more expressive, i.e., $\Sigma_{\sigma}^{ABA} \subseteq \Sigma_{\sigma_c}^{ABA}$. It turns out, however, that in general the opposite is the case. Let us start with the simple case of gr, where both notions indeed coincide (due to the simplicity of gr).

Proposition 3. $\Sigma_{gr}^{ABA} = \Sigma_{gr_c}^{ABA}$.

Proof. Each extension-set of size 1 can be realized for gr semantics in ABA. Our above construction does not require any rules, thus $\sigma(\mathcal{D}) = \sigma_c(\mathcal{D})$ □

Now, in general it is the case that assumption-extensions are more flexible in their modeling capabilities in the sense that $\Sigma_{\sigma_c}^{ABA} \subseteq \Sigma_{\sigma}^{ABA}$. To achieve this result we require a detour to so-called claim-augmented AFs [19] and their relation to SETAFs [18] and ABAFs [24]. We omit the proof details.

Proposition 4. *For all semantics σ in this paper, $\Sigma_{\sigma_c}^{ABA} \subseteq \Sigma_{\sigma}^{ABA}$.*

Interestingly, the other direction fails. For all semantics except gr, there are sets which are realizable as assumption-extensions, but not as conclusion extensions.

Proposition 5. *For all semantics $\sigma \neq gr$ in this paper, $\Sigma_{\sigma_c}^{ABA} \subsetneq \Sigma_{\sigma}^{ABA}$.*

Proof. (stb and pr) Let $\mathbb{S} = \{\{a\}, \{b\}, \{c\}\}$. We remark that $\{a\}, \{b\}, \{c\}$ can of course be realized by means of the usual assumption-extensions $\sigma(\mathcal{D})$ for pr and stb semantics, as our signatures results show. Now suppose $\sigma_c(\mathcal{D}) = \{\{a\}, \{b\}, \{c\}\}$.

We first argue that each of the three elements a, b, and c has to be an assumption: Supposing, e.g., that $a \notin \mathcal{A}$ holds yields that $a \in Th_{\mathcal{D}}(\emptyset)$, because otherwise $\{a\}$ could not be a conclusion-extension. However, in this case a would occur in each extension, but it does not. By symmetry, $\{a, b, c\} \subseteq \mathcal{A}$ must hold.

Thus $Th_{\mathcal{D}}(\{a\}) = \{a\}$, $Th_{\mathcal{D}}(\{b\}) = \{b\}$, and $Th_{\mathcal{D}}(\{c\}) = \{c\}$ for otherwise the conclusion-extensions would be larger. Since $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$ are no extension, we deduce that these sets are not conflict-free or not capable of defending themselves. However, the latter can be excluded since each single assumption a, b, c defends itself. So they have to be conflicting.

We distinguish several cases. (1) Suppose $\overline{a} = b$. But then a can only defend itself if $\overline{b} = a$ holds as well. Since $\{a, c\}$ is conflicting, it must also be the case that (a) $\overline{c} = a$ or (b) $\overline{a} = c$. Supposing (a), then $\{c\}$ can only defend itself if $\overline{a} = c$, i.e., this case yields (b). Supposing (b) implies $b = c$ since $^-$ is a function. But $b = c$ implies the extensions are actually $\{\{a\}, \{b\}\}$, a contradiction.

Other cases like, e.g., $\overline{b} = a$ yield analogous contradictions.

(ad and co) Now consider $\mathbb{S} = \{\emptyset, \{a\}, \{b\}, \{c\}\}$. Now, $\mathbb{S}$ can be realized by assumption-extensions *w.r.t.* ad and co semantics. Regarding σ_c, the same reasoning as above applies: $Th_{\mathcal{D}}(\{a\}) = \{a\}$, $Th_{\mathcal{D}}(\{b\}) = \{b\}$, and $Th_{\mathcal{D}}(\{c\}) = \{c\}$ can be inferred analogously and consequently, we find again that, e.g., $\{a\}$ cannot defend itself, yielding the same contradiction. □

Signatures for ABA with Preferences

Let us now head back to assumption-extensions. ABA^+ has been introduced in [14]; it generalizes ABA by incorporating preferences between the assumptions. We recall the necessary background.

Definition 13. *An* ABA^+ *framework is a tuple* $(\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-, \leq)$, *where* $(\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ *is an ABAF and* $\leq$ *is a transitive binary relation on* $\mathcal{A}$.

As usual, we write $a < b$ if $a \leq b$ and $b \not\leq a$. Attacks are generalized as follows.

Definition 14. *Given an* ABA^+ *framework* $(\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-, \leq)$. *A set of assumptions* $S \subseteq \mathcal{A}$ attacks *a set of assumptions* $T \subseteq \mathcal{A}$ *iff*

- *there is* $S' \subseteq S$, $t \in T$ *s.t.* $S' \vdash \overline{t}$, *and there is no* $s \in S'$ *with* $s < t$; *or*
- *there is* $T' \subseteq T$, $s \in S$ *s.t.* $T' \vdash \overline{s}$, *and there is* $t \in T'$ *with* $t < s$.

For ABA without preferences, only the first item matters: a set of assumptions S attacks another set of assumptions T iff (a subset of) S derives a contrary of some assumption in T. Taking preferences into account might cause an attack reversal, as formalized in item two. The semantics are defined as in Definition 2, but with the generalized attack notion as stated above. That is, S is *admissible* iff i) S does not attack itself and ii) if T attacks S, then S attacks T as well; S is *complete* iff it also contains each a it defends; S is *grounded* iff it is minimal complete and S *preferred* iff it is maximal admissible; S is stable iff S attacks each singleton $\{a\} \subseteq \mathcal{A} \backslash S$.

We let $\Sigma_\sigma^{ABA^+}$ denote the signature of semantics σ for ABA^+, i.e.,

$$\Sigma_\sigma^{ABA^+} = \{\sigma(\mathcal{D}) \mid \mathcal{D} \text{ is a flat, finite ABA}^+ \text{ framework}\}.$$

In this section, we establish the following main theorem.

Theorem 3. *It holds that*

$$\Sigma_{gr}^{ABA^+} = \{\mathbb{S} \mid |\mathbb{S}| \leq 1\},$$
$$\Sigma_{ad}^{ABA^+} = \{\mathbb{S} \mid \emptyset \in \mathbb{S}\},$$
$$\Sigma_{stb}^{ABA^+} = \{\mathbb{S} \mid \mathbb{S} \text{ is incomparable}\}, \text{ and}$$
$$\Sigma_{pr}^{ABA^+} = \Sigma_{stb}^{ABA} \setminus \{\emptyset\}.$$

We first note that each extension-set $\mathbb{S}$ which is contained in Σ_σ^{ABA} is also contained in $\Sigma_\sigma^{ABA^+}$; it suffices to consider the empty preference relation.

Proposition 6. *For all semantics considered in this paper* $\Sigma_\sigma^{ABA^+} \supseteq \Sigma_\sigma^{ABA}$.

For preferred and stable semantics, we even have a stronger result. Below, we show that the signatures for ABA$^+$ corresponds to the signatures for ABA without preferences. For preferred semantics, we obtain this result because the semantics operate on maximizing the assumption-sets. For stable semantics, we additionally rely on the monotonicity of the range function.

Proposition 7. *For* $\sigma \in \{stb, pr\}$ *we have* $\Sigma_\sigma^{ABA^+} = \Sigma_\sigma^{ABA}$.

Proof. Let $\sigma \in \{stb, pr\}$. According to Proposition 6, $\Sigma_\sigma^{ABA^+} \supseteq \Sigma_\sigma^{ABA}$ holds. On the other hand, $\Sigma_{pr}^{ABA^+} \subseteq \Sigma_{pr}^{ABA}$ since preferred extensions are incomparable by definition. For $\Sigma_{stb}^{ABA^+} \subseteq \Sigma_{stb}^{ABA}$ suppose $\mathcal{D}$ is an ABA^+ framework with stable extensions $S \subsetneq S'$. Even with preferences, the range is monotonic, i.e., we have that $Th_\mathcal{D}(S) \subseteq Th_\mathcal{D}(S')$. Consequently, if S is stable, then S' is not conflict-free; contradiction. □

In contrast, admissible semantics in ABA$^+$ are significantly more powerful than their counterpart in ABA without preferences. We observe that each extension set that contains the empty set can be realized.

Proposition 8. $\Sigma_{ad}^{ABA^+} = \{\mathbb{S} \mid \emptyset \in \mathbb{S}\}$.

Proof. Let $\mathbb{S} \neq \emptyset$ with $\emptyset \in \mathbb{S}$. Moreover, let $A_\mathbb{S} = \bigcup_{S \in \mathbb{S}} S$ and let $N_\mathbb{S} = 2^{A_\mathbb{S}} \setminus \mathbb{S}$. We construct the corresponding ABAF $\mathcal{D} = (\mathcal{L}, \mathcal{A}, \mathcal{R}, \overline{\ }, \leq)$ with

$$\mathcal{L} = \mathcal{A} \cup \{a^c \mid a \in \mathcal{A}\},\ \mathcal{A} = A_\mathbb{S} \cup \{v_N \mid N \in \mathbb{N}_\mathbb{S}\},\ \overline{a} = a^c \text{ for each } a \in \mathcal{A},$$
$$\mathcal{R} = \{{v_N}^c \leftarrow N; {v_N}^c \leftarrow v_N \mid N \in N_\mathbb{S}\} \cup \{{v_N}^c \leftarrow S \setminus N \mid S \in \mathbb{S}, N \in N_\mathbb{S}, N \subseteq S\},$$

and preference relation as follows: for each $N \in N_\mathbb{S}$, we let $v_N > n$ for some $n \in N$. An example of the construction is given in Fig. 1.

We observe that the set $A_\mathbb{S}$ itself is conflict-free(no assumption is a contrary). Each set N that is *not* contained in $\mathbb{S}$ receives an attack from v_N: The attack from the rule ${v_N}^c \leftarrow N$ is reversed because N contains some $n \in N$ with $n < v_N$. However, each set $S \in \mathbb{S}$ that is attacked by some v_N defends itself: v_N is counter-attacked by $S \setminus N$ (and there is no $s < v_N$ for any $s \in S \setminus N$ since $(S \setminus N) \cap N = \emptyset$). Based on these observations, we can show that $ad(\mathcal{D}) = \mathbb{S}$.

$$\begin{aligned}
\mathcal{L} &= \mathcal{A} \cup \{a^c \mid a \in \mathcal{A}\}\\
\mathcal{A} &= \{a, b, c, v_{\{b\}}, v_{\{c\}}, v_{\{a,b,c\}}\}\\
\mathcal{R} &= \{{v_{\{b\}}}^c \leftarrow b; \quad {v_{\{c\}}}^c \leftarrow c;\\
&\quad {v_{\{a,b,c\}}}^c \leftarrow a, b, c;\\
&\quad {v_{\{b\}}}^c \leftarrow a; \quad {v_{\{b\}}}^c \leftarrow c;\\
&\quad {v_{\{c\}}}^c \leftarrow a; \quad {v_{\{c\}}}^c \leftarrow b;\\
&\quad {v_{\{b\}}}^c \leftarrow v_{\{b\}}; \quad {v_{\{c\}}}^c \leftarrow v_{\{c\}};\\
&\quad {v_{\{a,b,c\}}}^c \leftarrow v_{\{a,b,c\}}\},
\end{aligned}$$

$$v_{\{b\}} > b, v_{\{c\}} > c, v_{\{a,b,c\}} > a$$

(a) Resulting ABA$^+$.

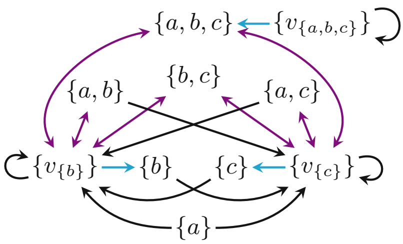

(b) Attacks between assumption-sets.

Fig. 1. Example of the construction from the proof of Proposition 8 for $\mathbb{S} = \{\emptyset, \{a, b\}, \{a, c\}, \{b, c\}, \{a\}\}$. We get $A_{\mathbb{S}} = \{a, b, c\}$ and $N_{\mathbb{S}} = \{\{b\}, \{c\}, \{a, b, c\}\}$. The corresponding ABA$^+$ is depicted left (1a), sets between assumption-sets are depicted right (1b) (supersets of $\{v_N\}$ are omitted since they are self-attacking).

First, we show that $ad(\mathcal{D}) \subseteq \mathbb{S}$. We note that $\emptyset \in ad(\mathcal{D})$ by definition. Now let $S \in \mathbb{S}$ with $S \neq \emptyset$. As observed above, $S \in cf(\mathcal{D})$. We show that S defends itself: let $X \subseteq \mathcal{A}$ be an assumption-set that attacks S. That is, either (a) there is $X' \subseteq X$, $s \in S$ such that $X' \vdash s$ and none of the elements in X' is strictly weaker than s, or (b) there is $S' \subseteq S$, $x \in X$ such that $S' \vdash x$ and there is $s \in S'$ such that $s < x$. By construction, no contrary of an assumption $a \in A_{\mathbb{S}}$ can be derived. Hence, case (a) cannot occur. Now, suppose (b) is the case. Then there are $S' \subseteq S$ and $x \in X$ such that $S' \vdash x$ and there is $s \in S'$ such that $s < x$. By construction, this can only be the case if S' is not contained in $\mathbb{S}$ (i.e., $S' \in N_{\mathbb{S}}$), S' derives the contrary of $v_{S'}$, and the direction of the attack is reversed by the preference relation $v_{S'} > n$ for some $n \in S'$. It is clear that S' is a *proper* subset of S. Otherwise we obtain $S \in N_{\mathbb{S}}$, contradicting our assumption. Hence, the set $S \setminus S'$ is not empty. By construction, the assumption $v_{S'}$ is attacked by $S \setminus S'$ (via the rule $v_{S'} \leftarrow S \setminus S'$). Hence, we obtain that S defends itself against the attack from X, as desired.

It remains to show that no other set is admissible, i.e., $ad(\mathcal{D}) \supseteq \mathbb{S}$. This is ensured by construction since each $N \in N_{\mathbb{S}}$ is attacked by v_N and is not defended against this assumption. We obtain $ad(\mathcal{D}) = \mathbb{S}$, as desired. □

Regarding grounded semantics, we require some auxiliary observations. We define the characteristic function for an ABA$^+$ framework $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-, \leq)$ as usual, i.e., we let $\Gamma_{\mathcal{D}}(S) = \{a \in \mathcal{A} \mid S \text{ defends } a\}$. The more involved attack notion does not alter the fact that Γ is monotonic.

Proposition 9. *Let $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-, \leq)$ be an ABA$^+$ framework. If $S \subseteq S' \subseteq \mathcal{A}$, then $\Gamma_{\mathcal{D}}(S) \subseteq \Gamma_{\mathcal{D}}(S')$.*

Monotonicity of the characteristic function is one of the key ingredients for showing that the grounded extension is unique (for the most common argumentation formalisms). Consequently, we can infer a similar result: The only candidate for the grounded extension is $\bigcup_{i\geq 1} \Gamma^i_{\mathcal{D}}(\emptyset)$, i.e., iterating the characteristic function. Perhaps somewhat surprising we can only derive $|gr(\mathcal{D})| \leq 1$, since complete extensions do not necessarily exist.

Proposition 10. *For any ABA*$^+$ *framework* $\mathcal{D} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-, \leq)$, $|gr(\mathcal{D})| \leq 1$.

As the following example shows, the case $co(\mathcal{D}) = \emptyset$ (and thus $gr(\mathcal{D}) = \emptyset$) is indeed possible. Consider the following simple ABA$^+$ framework.

Example 8. Let $\mathcal{D}$ be the ABA$^+$ framework with $\mathcal{A} = \{a, b, c\}$, the rule $\overline{c} \leftarrow a, b$, and let $c > a$. Then $\{c\}$ attacks $\{a, b\}$ since $\{a, b\} \vdash \overline{c}$ and $a < c$.

In $\mathcal{D}$, all assumptions a, b, and c are unattacked; however, the set $\{a, b, c\}$ is conflicting. Hence, no grounded extension exist in $\mathcal{D}$. Therefore, also complete semantics return the empty set. $\lozenge$

Consequently, the grounded ABA$^+$ signature is given as follows.

Proposition 11. $\Sigma^{ABA^+}_{gr} = \{\mathbb{S} \mid |\mathbb{S}| \leq 1\}$

Thereby, the above examples shows how to realize $\mathbb{S} = \emptyset$, and if $|\mathbb{S}| = 1$, then the construction given for usual ABAFs suffices. From the propositions we inferred within this section, the desired Theorem 3 follows.

6 Conclusion

In this paper, we investigated several aspects of ABA expressiveness. We characterized the signatures of ABA semantics by connecting two recent developments in the field of formal argumentation: we used the close relation to SETAFs presented in [24] in order to benefit from the established SETAF signatures [17]. We amplified our investigation with several aspects which are central for understanding the expressiveness of ABA. In particular, we discussed the relation to conclusion-based ABA semantics and signatures for ABA with preferences.

Our notion of signatures is inspired by research on expressiveness in abstract argumentation formalisms [16,17,20]. We are not aware of any comprehensive investigation of signatures in structured argumentation in the literature.

Searching for Suitable Translations. Semantics-preserving translations between non-monotonic reasoning formalisms are well-studied [10,11,24,29]. They are useful for several reasons. First, they enable access to solvers and other tools that have been developed for the target formalism (see e.g., [32]). Second, translations from structured to abstract argumentation formalisms have gained increasing attention in the context of explainability. Abstract graph-based representations are intuitive and easy-to-understand by design; moreover, they are central for extracting argumentative explanations. Since AFs are particularly well-studied, they are oftentimes considered as the default target formalism. However, many translations to AFs often require auxiliary arguments which may cause an exponential blowup; moreover, they often preserve semantics only under projection.

The underlying issue becomes clear when looking at the signatures in the different formalisms: it turns out that many argumentation formalisms are more expressive than AFs [17,20]. In particular, our results show that flat ABAFs are closer in their nature to SETAFs than to Dung's AFs. Moreover, we show that the more advanced ABA fragments that we consider admit a higher expressiveness than flat ABA for most of the semantics. Building upon our insights, we identify the search for suitable abstract formalisms with matching expressiveness that capture ABA with preferences, conclusion-semantics of ABA, or even more general fragments like non-flat ABA as a challenging future work direction. Generally speaking, it would be interesting to put more emphasis on abstract formalisms with higher expressiveness, e.g., in order to obtain competitive instantiation-based ABA solvers or to extract argumentative explanations.

The Role of Signatures in Dynamics in Argumentation. In particular in order to push forward dynamics research in structured argumentation, understanding the modeling capabilities of a formalism is crucial: oftentimes dynamics research is driven by a certain goal like, e.g., enforcing a target set of conclusions or forgetting given elements of a knowledge base [1,2,5,7,9,25].

Example 9. Suppose we want to develop a forgetting operator that removes from each extension the element that should be forgotten. This notion is known as *persistence*. Our signature results indicate whether it is possible so satisfy such constraints. For instance, it becomes clear that we cannot construct a forgetting operator that satisfies persistence for stable semantics: for an ABAF $\mathcal{D}$ with $stb(\mathcal{D}) = \{\{a, b\}, \{b, c\}\}$, we run into an issue if we aim to forget the assumption a. The set $\{\{b\}, \{b, c\}\}$ is not incomparable and hence it cannot be realized, as we have shown.[4] $\diamond$

Recent studies on dynamics in structured argumentation show that we cannot rely on the corresponding results for AFs [30,31]. Hence, our results provide a solid theoretical foundation in order to understand what can be attained and what not. Moreover, understanding expressiveness is indispensable in order to extend this line of research to further dynamic tasks like belief revision [3].

Open Problems. The present work characterizes the expressiveness of ABA semantics in several aspects. Nonetheless, certain questions in this context remain open: i) precise signature characterizations for admissible, complete, preferred, and stable semantics for conclusion-based ABA semantics; ii) precise signature characterizations of complete semantics for ABA^{+}; and iii) signature characterizations for semantics in non-flat ABAFs. We view closing these gaps as a natural future work directions. Moreover, our research was focusing on ABAFs, but there a several other structured argumentation formalisms worth investigating, for instance defeasible logic programming [23] or ASPIC+ [26].

[4] We refer the interested reader to [7] for an in-depth study on forgetting in flat ABA.

Acknowledgements. This research has been supported 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. Anna Rapberger was partially funded by the Vienna Science and Technology Fund (WWTF) through project ICT19-065, by the Austrian Science Fund (FWF) through project P32830, and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101020934).

References

1. Baumann, R.: What does it take to enforce an argument? Minimal change in abstract argumentation. In: Proceeding of (ECAI-12), pp. 127–132 (2012)
2. Baumann, R., Berthold, M.: Limits and possibilities of forgetting in abstract argumentation. In: Proceedings of (IJCAI-22), pp. 2539–2545. ijcai.org (2022)
3. Baumann, R., Brewka, G.: AGM meets abstract argumentation: Expansion and revision for dung frameworks. In: Proceedings of (IJCAI-15), pp. 2734–2740 (2015)
4. Baumann, R., Dvorák, W., Linsbichler, T., Strass, H., Woltran, S.: Compact argumentation frameworks. In: Proceedings of (ECAI-14). FAIA, vol. 263, pp. 69–74. IOS Press (2014)
5. Baumann, R., Gabbay, D.M., Rodrigues, O.: Forgetting an argument. In: Proceedings of (AAAI-20), pp. 2750–2757. AAAI Press (2020)
6. Baumann, R., Strass, H.: On the maximal and average numbers of stable extensions. In: Black, E., Modgil, S., Oren, N. (eds.) TAFA 2013. LNCS (LNAI), vol. 8306, pp. 111–126. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54373-9_8
7. Berthold, M., Rapberger, A., Ulbricht, M.: Forgetting aspects in assumption-based argumentation. In: Proceedings of (KR-23) (2023)
8. Bondarenko, A., Toni, F., Kowalski, R.A.: An assumption-based framework for non-monotonic reasoning. In: Proceedings of (LPNMR-93), pp. 171–189. MIT Press (1993)
9. Borg, A., Bex, F.: Enforcing sets of formulas in structured argumentation. In: Proceedings of (KR-21), pp. 130–140 (2021)
10. Caminada, M., Sá, S., Alcântara, J., Dvořák, W.: On the difference between assumption-based argumentation and abstract argumentation. IFCoLog J. Logic Appl. **2**(1), 15–34 (2015)
11. Caminada, M., Sá, S., Alcântara, J., Dvořák, W.: On the equivalence between logic programming semantics and argumentation semantics. Int. J. Approximate Reasoning **58**, 87–111 (2015)
12. Čyras, K., Fan, X., Schulz, C., Toni, F.: Assumption-based argumentation: disputes, explanations, preferences. In: Handbook of Formal Argumentation, vol. 1, chap. 7, pp. 365–408. College Publications (2018)
13. Cyras, K., Rago, A., Albini, E., Baroni, P., Toni, F.: Argumentative XAI: a survey. In: Proceedings of (IJCAI-21), pp. 4392–4399. ijcai.org (2021)
14. Cyras, K., Toni, F.: ABA+: assumption-based argumentation with preferences. CoRR abs/1610.03024 (2016)
15. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artif. Intell. **77**(2), 321–357 (1995)

16. Dunne, P.E., Dvořák, W., Linsbichler, T., Woltran, S.: Characteristics of multiple viewpoints in abstract argumentation. Artif. Intell. **228**, 153–178 (2015)
17. Dvorák, W., Fandinno, J., Woltran, S.: On the expressive power of collective attacks. Argument Comput. **10**(2), 191–230 (2019)
18. Dvorák, W., Rapberger, A., Woltran, S.: On the relation between claim-augmented argumentation frameworks and collective attacks. In: Proceedings of (ECAI-20). FAIA, vol. 325, pp. 721–728. IOS Press (2020)
19. Dvorák, W., Woltran, S.: Complexity of abstract argumentation under a claim-centric view. Artif. Intell. **285**, 103290 (2020)
20. Dvořák, W., Rapberger, A., Woltran, S.: Argumentation semantics under a claim-centric view: Properties, expressiveness and relation to SETAFs. In: Proceedings of (KR-20), pp. 341–350 (2020)
21. Fan, X., Toni, F.: On computing explanations in argumentation. In: Proceedings of (AAAI-15), pp. 1496–1502. AAAI Press (2015)
22. Gabbay, D., Giacomin, M., Simari, G.R., Thimm, M. (eds.): Handbook of Formal Argumentation, vol. 2. College Publications (2021)
23. García, A.J., Simari, G.R.: Defeasible logic programming: an argumentative approach. Theory Pract. Logic Program. **4**(1–2), 95–138 (2004)
24. König, M., Rapberger, A., Ulbricht, M.: Just a matter of perspective. In: Proceedings of (COMMA-22). FAIA, vol. 353, pp. 212–223. IOS Press (2022)
25. Lin, F., Reiter, R.: Forget it. In: Working Notes of AAAI Fall Symposium on Relevance, pp. 154–159 (1994)
26. Modgil, S., Prakken, H.: The ASPIC$^+$ framework for structured argumentation: a tutorial. Argument Comput. **5**(1), 31–62 (2014)
27. Nielsen, S.H., Parsons, S.: A generalization of dung's abstract framework for argumentation: arguing with sets of attacking arguments. In: Maudet, N., Parsons, S., Rahwan, I. (eds.) ArgMAS 2006. LNCS (LNAI), vol. 4766, pp. 54–73. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75526-5_4
28. Odekerken, D., Borg, A., Berthold, M.: Accessible algorithms for applied argumentation. In: Proceedings of (Arg&App@KR-23) (2023)
29. Polberg, S.: Developing the Abstract Dialectical Framework. Phd thesis, Vienna University of Technology, Institute of Information Systems (2017)
30. Prakken, H.: Formalising an aspect of argument strength: Degrees of attackability. In: Proceedings of (COMMA-22), vol. 353, pp. 296–307. IOS Press (2022)
31. Rapberger, A., Ulbricht, M.: On dynamics in structured argumentation formalisms. J. Artif. Intell. Res. **77**, 563–643 (2023)
32. Tuomo, L., Rapberger, A., Ulbricht, M., Wallner, J.P.: Argumentation frameworks induced by assumption-based argumentation: relating size and complexity. In: Proceedings of (KR-23) (2023)
33. Ulbricht, M.: On the maximal number of complete extensions in abstract argumentation frameworks. In: Proceedings of (KR-21), pp. 707–711 (2021)
34. Vassiliades, A., Bassiliades, N., Patkos, T.: Argumentation and explainable artificial intelligence: a survey. Knowl. Eng. Rev. **36**, e5 (2021)