

Investigating subclasses of abstract dialectical frameworks

Martin Diller^a, Atefeh Keshavarzi Zafarghandi^{b,*}, Thomas Linsbichler^c and Stefan Woltran^{c,**}

^a *Computational Logic Group, TU Dresden, Germany*

^b *Department of Artificial Intelligence, Bernoulli Institute, University of Groningen, The Netherlands*

^c *Institute of Logic and Computation, TU Wien, Austria*

Abstract. Abstract dialectical frameworks (ADFs) are generalizations of Dung argumentation frameworks where arbitrary relationships among arguments can be formalized. This additional expressibility comes with the price of higher computational complexity, thus an understanding of potentially easier subclasses is essential. Compared to Dung argumentation frameworks, where several subclasses such as acyclic and symmetric frameworks are well understood, there has been no in-depth analysis for ADFs in such direction yet (with the notable exception of bipolar ADFs). In this work, we introduce certain subclasses of ADFs and investigate their properties. In particular, we show that for acyclic ADFs, the different semantics coincide. On the other hand, we show that the concept of symmetry is less powerful for ADFs and further restrictions are required to achieve results that are similar to the known ones for Dung's frameworks. A particular such subclass (support-free symmetric ADFs) turns out to be closely related to argumentation frameworks with collective attacks (SETAFs); we investigate this relation in detail and obtain as a by-product that even for SETAFs symmetry is less powerful than for AFs. We also discuss the role of odd-length cycles in the subclasses we have introduced. Finally, we analyse the expressiveness of the ADF subclasses we introduce in terms of signatures.

Keywords: Abstract argumentation, abstract dialectical frameworks, acyclic frameworks, symmetric frameworks, expressiveness

1. Introduction

Since the landmark paper by Dung [13] has been published in 1995, abstract argumentation frameworks (AFs) have gained more and more significance in the AI domain. First of all, AFs have proven useful to capture the essence of different nonmonotonic formalisms. Moreover, AFs are nowadays an integral concept in several advanced argumentation-based formalisms in the sense that their semantics are defined based on a translation (typically called an instantiation) to Dung AFs. Finally, the relevance of AFs is witnessed by the *International Competition on Computational Models of Argumentation (IC-CMA)*, where systems for solving different problems on AFs compete on different tracks.¹

The fundamental of Dung is to abstract away from the content of particular arguments and to focus only on conflicts among arguments, where each argument is viewed as an atomic item. Therefore, the only information AFs take into account is whether an argument attacks another one or not. Semantics

* Corresponding author. E-mail: a.keshavarzi.zafarghandi@rug.nl.

** Author names are sorted in alphabetical order. This paper is based on the second author's master thesis (A. Keshavarzi Zafarghandi (2017)).

¹<http://argumentationcompetition.org>

single out coherent subsets of arguments which “fit” together, according to specific criteria [2]. More formally, an AF semantics takes an argumentation framework as input and produces as output a collection of sets of arguments, called extensions. Complexity of the reasoning problems that can be defined for the several semantics for AFs is well understood [17] and ranges from tractability up to the second level of the polynomial hierarchy. To this end, the analysis of restricted classes of AFs is of importance. In his paper, Dung already showed that the class of acyclic (also known as well-founded) AFs leads to a collapse of the different semantics. Further studies include symmetric AFs [10] and AFs under other graph-driven restrictions [14]. Symmetric AFs have been proven to satisfy the property of coherence (preferred and stable semantics coincide) and relatively-groundedness (the grounded extension is given by the intersection of the preferred extensions). Moreover, these restrictions make decision problems often easier from a complexity perspective. A fact that is particularly useful in connection with backdoor approaches [20] that utilize the distance to an easier fragment. This approach has, for instance, been realised in practice with the cegartix system [19].

Abstract dialectical frameworks (ADFs) are generalizations of Dung argumentation frameworks where arbitrary relationships among arguments can be formalized via propositional formulas which are attached to the arguments [5,8]. This allows to express notions of support, collective attacks, and even more complicated relations. Due to their flexibility in formalizing relations between arguments, ADFs have recently been used in several applications [1,9,26,30]. However, this additional expressibility comes with the price of higher computational complexity [35]. Specifically, reasoning in ADFs spans the first three (rather than the first two, as for AFs) levels of the polynomial hierarchy.

It is thus natural to investigate subclasses of ADFs. Compared to Dung argumentation frameworks, where subclasses like acyclic and symmetric AFs have been thoroughly studied, there has not been a systematic investigation of subclasses of ADFs yet. An exception is the class of bipolar ADFs [8] where the links between arguments are restricted to have either supporting or attacking nature. However, results about structural restrictions on ADFs where different semantics coincide are still lacking.

In this work, we aim to define several subclasses of ADFs and investigate how the restrictions we define influence the semantic evaluation of such ADFs. As a first class, we consider acyclic ADFs (i.e., the link-structure forms an acyclic graph) and show that – analogous to well-founded AFs – the main semantics, namely grounded, complete, preferred, and two-valued model/stable semantics, coincide for this class. We further investigate the concept of symmetric ADFs. In contrast to the case of AFs, we will see that properties as coherence and relatively-groundedness do not carry over and require further restrictions which leads us to the classes of acyclic support symmetric ADFs (ASSADFs) and support-free symmetric ADFs (SFSADFs). For both classes we show that they satisfy a weaker form of coherence. We also show that these two classes differ in the sense that odd-cycle free SFSADFs are coherent while odd-cycle free ASSADFs are not. As a second contribution, following the work of Dunne et al. [16], we investigate the expressiveness of our ADF subclasses in terms of signatures, i.e. the set of possible outcomes which can be achieved by ADFs (of a particular class) under the different semantics. We thus complement here results which have been obtained for general [29,32] and bipolar ADFs [25] and also compare our ADF subclasses to abstract argumentation frameworks in terms of expressibility.

Our results lead to the following implications. Firstly, studying subclasses of ADFs provides us with a better understanding of which structures are required to reveal particular behaviors of the different semantics. We thus further advance the theory of ADFs. Secondly, since other generalizations of Dung AFs can be seen as special case of ADFs, results on ADFs carry over to these special cases. We exemplify this aspect in the paper, by deriving new results for argumentation frameworks with collective attacks (SETAFs) [27] which have received increasing interest recently [18,21]. To the best of our knowledge

concepts like symmetric SETAFs have not been investigated yet, and we provide first results in this direction.

The paper is structured as follows: In Section 2, we first recall the relevant background of AFs, SETAFs, and ADFs. Then in Section 3 we introduce subclasses of ADFs and we investigate whether these subclasses fulfill the same properties of the similar subclasses in AFs. We discuss how our results can be related to SETAFs and investigate some properties for symmetric SETAFs. Also the role of odd-length cycles is addressed. In Section 4, expressiveness of the subclasses of ADFs introduced in the current work is studied. In particular, we show that the expressiveness of SFSADFs, ASSADFs and bipolar ADFs is equal for some of the semantics, but different for admissibility-based semantics.

A preliminary version of this article appeared in [11]. This extended version contains new technical results including investigations concerning coherence and relatively-groundedness for SFSADFs and symmetric SETAFs (Theorems 8 and 17); and results on the role of odd-cycles on coherence for subclasses of ADFs (Section 3.4). Also, the results on expressibility for SETAFs and SFSADFs (as well as for a further superclass of SFSADFs we introduce, that of SFADFs) in Section 4 are new.

2. Preliminaries

Let us start by recalling basic notions of Dung's argumentation frameworks (AFs) and their extension for modelling collective attacks (SETAFs). We then introduce the formalism that is the focus of this work, abstract dialectical frameworks (ADFs), and point out a few elemental relations between ADFs, SETAFs, and AFs.

2.1. Abstract argumentation frameworks

Definition 1 ([13]). An *abstract argumentation framework* (AF) is a pair (A, R) in which A is a set of arguments and $R \subseteq A \times A$ is a binary relation representing attacks among arguments.

Let $F = (A, R)$ be an AF. When used for modelling purposes it is usually assumed that the set of arguments A is finite. Also, for each $a, b \in A$, the relation $(a, b) \in R$ is used to represent that a is an argument attacking the argument b . We can then also say that an argument $b \in A$ is attacked by a set $S \subseteq A$ if there is a $a \in S$ that attacks b . An argument $a \in A$ is, on the other hand, *defended* by a set $S \subseteq A$ of arguments (alternatively, the argument is *acceptable* with respect to S) (in F) if for each argument $c \in A$, it holds that if $(c, a) \in R$ then there is a $s \in S$ such that $(s, c) \in R$. In addition, a subset S of A is called *conflict-free* in F , if it does not attack itself. I.e. there are no $a, b \in S$ such that $(a, b) \in R$.

In abstract argumentation, diverse criteria that have been formulated for conflict-resolution are encoded into so called semantics. In the case of AFs, semantics identify sets of arguments which it is reasonable to accept together (w.r.t. the semantics at hand). One standard way of defining the semantics for a given AF $F = (A, R)$ is via the so called *characteristic operator* Γ_F : a function mapping a set of arguments of F to the set of arguments defended by itself. Formally, let S be a subset of A , then $\Gamma_F(S) = \{a \in A \mid a \text{ is defended by } S \text{ in } F\}$. We compile the definitions of some of the main semantics of AFs given in terms of the characteristic operator in the following:

Definition 2. Let $F = (A, R)$ be an AF. A set S which is conflict-free in F is:

- *admissible* in F iff $S \subseteq \Gamma_F(S)$;
- *complete* in F iff $S = \Gamma_F(S)$;
- *grounded* in F iff S is the \subseteq -least fixed-point of Γ_F ;
- *preferred* in F iff S is \subseteq -maximal admissible (resp. complete) in F ;
- a *stable model* of F iff for all $a \in A \setminus S$, S attacks a .

In order to denote the collection of all conflict-free sets, (admissible sets, complete sets, grounded sets, preferred sets, stable models) of an AF F , we use the notation $\sigma(F)$ where $\sigma \in \{cf, adm, com, grd, prf, stb\}$. Such sets are usually called extensions; we will also often refer to them as *extension-sets* or simply σ -sets (of arguments) where σ is a semantics.

We recall that, for each AF, the grounded set is the unique subset-minimal complete set. Moreover, for a given AF a stable model might not exist, while the remaining semantics produce at least one set of arguments for any AF. Note also that we use the notion of a “stable model” rather than “stable extension” simpliciter (as is more common for AFs) just for using terminology compatible with that used for the semantics of ADFs (where stable models are a subclass of two-valued-models) which we introduce later-on in Definition 6.

We now turn to the generalization of AFs introduced by Nielson and Parsons, in which the attack relation is generalised to also include attacks from sets of two or more arguments [27].

Definition 3. A set argumentation framework (SETAF) is an ordered pair $F = (A, R)$, where A is a finite set of arguments and $R \subseteq (2^A \setminus \{\emptyset\}) \times A$ is the attack relation.

Given a SETAF (A, R) , we write $S \mapsto_R b$ if there is a set $S' \subseteq S$ attacking b , i.e. $(S', b) \in R$. We say that in this case also S attacks b . Moreover, we write $S' \mapsto_R S$ if $S' \mapsto_R b$ for some $b \in S$. We drop the subscript in \mapsto_R if the attack relation is clear from the context.

Notions of conflict and defense can be defined for SETAFs in analogy to these notions in the context of AFs. Given a SETAF $F = (A, R)$, a set $S \subseteq A$ is *conflicting* in F if $S \mapsto_R S$; $S \subseteq A$ is *conflict-free* in F , if S is not conflicting in F , i.e. if $S' \cup \{a\} \not\subseteq S$ for each $(S', a) \in R$. An argument $a \in A$ is *defended* (in F) by a set $S \subseteq A$ if for each $B \subseteq A$, such that $B \mapsto_R a$, also $S \mapsto_R B$. A set T of arguments is hence defended (in F) by S if each $a \in T$ is defended by S (in F).

The semantics of SETAFs can now also be defined similarly to AFs via a characteristic operator. With a slight abuse of notation (because of the use of the same notation for the characteristic operator), we thus define first of all also for a SETAF $F = (A, R)$, $\Gamma_F(S) = \{a \in A \mid a \text{ is defended by } S \text{ in } F\}$; here the notion of “defense” clearly being that defined for SETAFs. For completeness we detail the definitions of all semantics we consider in this work for SETAFs, although the definitions are exactly as those for AFs (modulo the use of the more general notions of attack and that the characteristic operator referenced therein is the characteristic operator defined for SETAFs):

Definition 4. Let $F = (A, R)$ be a SETAF. A set S which is conflict-free in F is

- *admissible* in F iff $S \subseteq \Gamma_F(S)$;
- *complete* in F iff $S = \Gamma_F(S)$;
- *grounded* in F iff S is the \subseteq -least fixed-point of Γ_F ;
- *preferred* in F iff S is \subseteq -maximal admissible (resp. complete) in F ;
- a *stable model* of F iff for all $a \in A \setminus S$, S attacks a .




Fig. 1. The SETAF of example 1.

We will use the same abbreviations for SETAFs as for AFs for denoting the sets of arguments obtained when applying the semantics on SETAFs. We also recall that several important properties of Dung AFs carry over to SETAFs; we refer to [21,27] for details.

In the following we provide an example of a SETAF and also illustrate the concept of extensions and semantics for SETAFs.

Example 1. The SETAF $F = (\{a, b, c\}, \{(\{a, b\}, c), (\{a, c\}, b)\})$ is depicted in Fig. 1. In F , $(\{a, b\}, c) \in R$ says that there is a joint attack from a and b to c , and $(\{a, c\}, b) \in R$ says that there is a joint attack from a and c to b . The former attack represents that neither a nor b are strong enough to attack c by themselves. The latter attack indicates that neither a nor c are strong enough to attack b by themselves. The conflict-free extensions of F are $cf(F) = \{\{\}, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}\}$, the admissible extensions $adm(F) = \{\{\}, \{a\}, \{a, b\}, \{a, c\}\}$, the complete extensions $com(F) = \{\{a\}, \{a, b\}, \{a, c\}\}$, the unique grounded extension $grd(F) = \{\{a\}\}$, and the preferred extensions $prf(F) = stb(F) = \{\{a, b\}, \{a, c\}\}$. Note that, for instance, $\{b, c\}$ is a conflict-free extension. However, it is not an admissible extension, since $\{b, c\} \not\subseteq (\Gamma_F(\{b, c\}) = \{\})$. Further, $\{a\}$ is an admissible and a complete extension, since $\Gamma_F(\{a\}) = \{a\}$. On the other hand $\{a\}$ is not a preferred extension because it is not a \subseteq -maximal admissible extension.

2.2. Abstract dialectical frameworks

We now proceed to define ADFs:

Definition 5 ([8]). An *abstract dialectical framework* (ADF) is a tuple $D = (S, L, C)$ where;

- S is a finite set of nodes (arguments, statements, positions),
- $L \subseteq S \times S$ is a set of links,
- $C = \{\varphi_s\}_{s \in S}$ is a collection of propositional formulas (acceptance conditions).

The variables occurring in each acceptance condition $\varphi_s \in C$ are taken from the *parents* of s , $par(s) = \{a \in S \mid (a, s) \in L\}$. Thus the arguments of an ADF act at the same time as nodes of the ADF and variables in the acceptance conditions. An argument $s \in S$ is called *initial* if $par(s) = \emptyset$.

Graphically, ADFs can be depicted as annotated directed graphs where nodes represent arguments, directed edges represent links (tuples of nodes). Acceptance conditions are shown via annotations next to the nodes representing the arguments. See Fig. 2 for an example of the depiction of an ADF (we present later on in Example 2).

Semantics of ADFs are defined in terms of interpretations. Concretely, let $D = (S, L, C)$ be an ADF. An *interpretation* v (for D) is then a function mapping each of the arguments in S to one of the truth

values true (**t**), false (**f**), or undecided (**u**), i.e. $v : S \mapsto \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$. The truth values stand for different degrees of acceptance that can be associated to the arguments. The sets $v^{\mathbf{t}}$, $v^{\mathbf{f}}$, and $v^{\mathbf{u}}$ contain those arguments that v maps to true, false and undecided, respectively. An interpretation v is *two-valued* if it maps each argument to either **t** or **f**, i.e. $v^{\mathbf{u}} = \emptyset$. At the other extreme, an interpretation v is called *trivial*, denoted $v_{\mathbf{u}}$, if $v(s) = \mathbf{u}$ for each $s \in S$. Also, we denote the update of an interpretation v with a truth value $\mathbf{x} \in \{\mathbf{t}, \mathbf{f}\}$ for an argument b by $v|_{\mathbf{x}}^b$, i.e. $v|_{\mathbf{x}}^b(b) = \mathbf{x}$ and $v|_{\mathbf{x}}^b(a) = v(a)$ for $a \neq b$.

Let \mathcal{V} be the set of all interpretations for an ADF D as before. Then, we call a subset of all interpretations of the ADF, $\mathbb{V} \subseteq 2^{\mathcal{V}}$, an *interpretation-set*. Interpretations can now be ordered with respect to their information content. This is based on the partial order of truth values, for which $\mathbf{u} \leq_i \mathbf{t}$ and $\mathbf{u} \leq_i \mathbf{f}$. An interpretation v is at least as informative as another interpretation w , denoted by $w \leq_i v$, if $w(s) \leq_i v(s)$ for each $s \in S$. As usual $w <_i v$ whenever $w \leq_i v$ but it is not the case that $v \leq_i w$. The *meet* operator \sqcap_i on truth values is defined as $\mathbf{t} \sqcap_i \mathbf{t} = \mathbf{t}$, $\mathbf{f} \sqcap_i \mathbf{f} = \mathbf{f}$ and returns \mathbf{u} otherwise. The meet \sqcap_i of two interpretations v and w is then defined as $(v \sqcap_i w)(s) = v(s) \sqcap_i w(s)$ for each $s \in S$. The partially ordered set $(\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}, \leq_i)$ forms a meet-semilattice with the meet operator \sqcap_i .

Semantics for ADFs can also be defined via a *characteristic operator* Γ_D for an ADF D , although given that the semantics of ADFs give interpretations rather than sets of arguments, the definitions now deviate somewhat to that for AFs and SETAFs. Given an interpretation v (for D), the characteristic operator Γ_D for D is defined as

$$\Gamma_D(v) = v' \quad \text{such that } v'(s) = \begin{cases} \mathbf{t} & \text{if } \varphi_s^v \text{ is irrefutable (i.e., a tautology),} \\ \mathbf{f} & \text{if } \varphi_s^v \text{ is unsatisfiable,} \\ \mathbf{u} & \text{otherwise,} \end{cases}$$

where the partial valuation of φ_s by v , is given by $\varphi_s^v = v(\varphi_s) = \varphi_s[p/\top : v(p) = \mathbf{t}][p/\perp : v(p) = \mathbf{f}]$. Here $p \in \text{par}(s)$. We will also sometimes write that $s \in S$ is *acceptable* w.r.t. v (in D) if φ_s^v is irrefutable. On the other hand $s \in S$ is *deniable* w.r.t. v (in D) if φ_s^v is unsatisfiable. Importantly, the characteristic operator for ADFs is monotonic: if $w \leq_i v$ for a given ADF D , then also $\Gamma_D(w) \leq_i \Gamma_D(v)$.

The semantics for ADFs, as defined via the characteristic operator, are provided next in Definition 6. The definitions have been introduced in [4,8]; we refer to [6] for an in-depth overview of the semantics of ADFs in terms of approximation fixpoint theory and the connections to the definitions of the semantics for AFs (and SETAFs).

Definition 6. Given an ADF $D = (S, L, C)$, an interpretation v is

- *conflict-free* in D iff $v(s) = \mathbf{t}$ implies φ_s^v is satisfiable and $v(s) = \mathbf{f}$ implies φ_s^v is unsatisfiable;
- *admissible* in D iff $v \leq_i \Gamma_D(v)$;
- *complete* in D iff $v = \Gamma_D(v)$;
- *grounded* in D iff v is the least fixed-point of Γ_D ;
- *preferred* in D iff v is \leq_i -maximal admissible (resp. complete) in D ;
- a *(two-valued) model* of D iff v is two-valued and for all $s \in S$, it holds that $v(s) = v(\varphi_s)$;
- a *stable model* of D if v is a model of D and $v^{\mathbf{t}} = w^{\mathbf{t}}$, where w is the grounded interpretation of the *stb*-reduct $D^v = (S^v, L^v, C^v)$, where $S^v = v^{\mathbf{t}}$, $L^v = L \cap (S^v \times S^v)$, and $\varphi_s[p/\perp : v(p) = \mathbf{f}]$ for each $s \in S^v$.

As for AFs and SETAFs, the set of all σ interpretations for D is denoted by $\sigma(D)$, where $\sigma \in \{\text{cf}, \text{adm}, \text{com}, \text{grd}, \text{prf}, \text{mod}, \text{stb}\}$ abbreviates the different semantics in the obvious manner.

In ADFs links between arguments or statements can be classified into four types, reflecting the relationship of attack and/or support that exists between the statements. These are listed in the following definition.

Definition 7. Let $D = (S, L, C)$ be an ADF. A link $(b, a) \in L$ is called

- *supporting* (in D) if for every two-valued interpretation v , $v(\varphi_a) = \mathbf{t}$ implies $v|_{\mathbf{t}}^b(\varphi_a) = \mathbf{t}$;
- *attacking* (in D) if for every two-valued interpretation v , $v(\varphi_a) = \mathbf{f}$ implies $v|_{\mathbf{t}}^b(\varphi_a) = \mathbf{f}$;
- *redundant* (in D) if it is both attacking and supporting;
- *dependent* (in D) if it is neither attacking nor supporting.

The classification of the types of the links of ADFs is also relevant for classifying ADFs themselves. Thus, one particularly important subclass of ADFs is that of *bipolar* ADFs or BADFs for short, first defined in [8]. In such an ADF each link is either attacking or supporting (or both; thus, the links can also be redundant). The following example clarifies the role of the acceptance conditions in ADFs and illustrates the functioning of the different semantics in the context of ADFs.

Example 2. An example of an ADF $D = (S, L, C)$ is shown in Fig. 2. To each argument a propositional formula is associated, the acceptance condition of the argument. In the acceptance condition of an argument only the parents of that argument occur as variables. The acceptance condition indicates how the acceptance-status of an argument depends on that of its parents in the ADF. For instance, the acceptance condition of c , namely $\varphi_c : \neg a \vee \neg b$, states that c can be accepted in an interpretation where either a or b (or both) are rejected. The acceptance condition $\varphi_a : \perp$ states that a always has to be rejected. The acceptance condition $\varphi_b : a \vee b \vee \neg c$ indicates that b can be accepted if and only if a or b is accepted or c is rejected, thus also indicating a form of self-support for b .

Since in ADFs an argument appears in the acceptance condition of an argument a if and only if it belongs to the set $par(a)$, the set of links L of an ADF is given implicitly via the acceptance conditions. Thus we often also don't give an explicit representation of the links in the examples we provide in this work. For instance, since b appears in φ_b , $(b, b) \in L$. Moreover, given that for every two-valued interpretation v , $v(\varphi_b) = \mathbf{t}$ implies $v|_{\mathbf{t}}^b(\varphi_b) = \mathbf{t}$, the link (b, b) is supporting (in D). Further, D is a bipolar ADF in which $L^+ = \{(a, b), (b, b)\}$ is the set of all supporting links, and $L^- = \{(a, c), (b, c), (c, b)\}$ is the set of all attacking links.

In D the interpretation $v = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{t}\}$ is conflict-free. However, v is not an admissible interpretation, because $\Gamma_D(v) = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}\}$, that is, $v \not\leq_i \Gamma_D(v)$. The interpretation $v_1 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}\}$ on the other hand is an admissible interpretation and, in fact, also the unique grounded interpretation of D . Further, $prf(D) = mod(D) = \{\{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}\}\}$, but only the first interpretation in this set is a stable model. This is because for $v = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}$ the unique grounded interpretation w of D^v is $\{c \mapsto \mathbf{t}\}$ and $v^{\mathbf{t}} = w^{\mathbf{t}}$.




Fig. 2. The ADF of example 2.

The interpretation $v' = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}\}$ is not a stable model, since the unique grounded interpretation w' of $D^{v'}$ is $\{b \mapsto \mathbf{u}, c \mapsto \mathbf{t}\}$ and $v'^{\mathbf{t}} \neq w'^{\mathbf{t}}$. In addition, we have that for the ADF D , $\text{com}(D) = \text{prf}(D) \cup \text{grd}(D)$.

As to the relationship between ADFs and AFs, an ADF $D = (S, L, C)$ is AF-like if the acceptance condition of every argument $s \in S$ is of the form $\varphi_s = \bigwedge_{a \in \text{par}(s)} \neg a$ (by the usual conventions, we thus have that if $\text{par}(s) = \emptyset$, then $\varphi_s = \top$). Such an ADF can be represented as an AF, as a tuple $F = (A, R)$ with arguments $A = S$ and attacks $R = L$. In turn, each AF can be represented as an ADF via the following definition.

Definition 8. Let $F = (A, R)$ be an AF. The associated ADF of F is $D_F = (S, L, C)$ in which $S = A$, $L = R$, and the acceptance condition of each $s \in S$ in C is given as $\varphi_s = \bigwedge_{a:(a,s) \in R} \neg a$.

It is shown in [4,8] that the semantics of AFs and associated ADFs coincide. Also note that the associated ADF D_F of each AF F is a BADF in which all links are attacking. In fact, BADFs are a strict superclass of AFs and strict subclass of ADFs.

An ADF $D = (S, L, C)$ is, on the other hand, SETAF-like if the acceptance condition of every argument $s \in S$ is given by a formula where each occurrence of an atom has negative polarity and is different to \perp . The reason behind this definition is the fact that in [36] it is shown that an ADF can directly be written as a SETAF if the acceptance condition of each argument is in conjunctive normal form and consists only of negative literals (note that arbitrary formulas where each atom occurs only having negative polarity can be transformed to such a CNF). We come back to this issue later in Section 3.3 where we introduce a further subclass of ADFs, support-free ADFs, that contains the class of SETAFs.

To conclude the preliminaries to our work we note that, just as AFs, also SETAFs can be represented as ADFs. This is captured in the following definition.

Definition 9. Let $F = (A, R)$ be a SETAF. The ADF associated to F is a tuple $D_F = (S, L, C)$ in which, $S = A$, $L = R$ and $C = \{\varphi_a\}_{a \in S}$ is the collection of acceptance conditions defined, for each $a \in S$, as

$$\varphi_a = \bigwedge_{(B,a) \in R} \bigvee_{a' \in B} \neg a'.$$

3. Properties of ADF subclasses

We start our investigation of ADF subclasses in terms of their semantics by first introducing the class of acyclic ADFs and showing that, just as is the case for acyclic AFs [13], the different semantics coincide on such ADFs. Then, we consider symmetric ADFs, where we will explore further restrictions that are needed in order to obtain results similar to the ones known for symmetric AFs. In Section 3.3 we discuss implications of our results for SETAFs. We conclude this section with a brief overview on semantic properties of odd-cycle free subclasses of ADFs.

3.1. Acyclic ADFs

In this section we show that, as has already been indicated and is the case for acyclic AFs, also for acyclic ADFs several semantics coincide (Theorem 2). We start by defining acyclic ADFs.

Definition 10. An ADF $D = (S, L, C)$ is *acyclic* if its corresponding directed graph (S, L) is acyclic.

In order to prove that the different semantics coincide on acyclic ADFs we need the concepts of level and maximum level of arguments. The *level* of an argument s of an ADF D is the number of links on the longest path from an initial argument to s plus 1. The *maximum level* of an (acyclic) ADF D then is the level of an argument of D that is at least as large as the level of any other argument of D . It is clear that every acyclic ADF has a maximum level. This is a crucial observation needed for our proof of Proposition 1, which in turn provides the basis to show that most semantics defined for ADFs are indistinguishable when evaluating acyclic ADFs.

Proposition 1. *In every acyclic ADF D the \leq_i -least fixed point of Γ_D is a model of D .*

Proof. Let $D = (S, L, C)$ be an acyclic ADF and let m be its maximum level. Moreover, let $v_0 := v_{\mathbf{u}}$ and $v_i := \Gamma_D(v_{i-1})$ for $1 \leq i \leq m$. We claim that for all i with $1 \leq i \leq m$, and every argument s_j with level $j \leq i$ it holds that either $v_i(s_j) = \mathbf{t}$ or $v_i(s_j) = \mathbf{f}$. We show this claim by induction on i :

- Base case: Suppose s_1 is an arbitrary argument of level one (an acyclic ADF always includes an initial argument). Since s_1 is an initial argument, either $\varphi_{s_1} = \top$ or $\varphi_{s_1} = \perp$. Hence $v_1(s_1) = \Gamma_D(v_0)(s_1)$ is either true or false.
- Inductive step: Assuming this property holds for all k with $1 \leq k \leq i < m$, we show it holds for $i + 1$. We know that $\varphi_{s_j}^{v_i} = \varphi_{s_j}[s_k/\top : v_i(s_k) = \mathbf{t}][s_k/\perp : v_i(s_k) = \mathbf{f}]$. For all s_k that occur in φ_{s_j} it holds that $k < j \leq i + 1$, with k being the level of s_k . Therefore, by the inductive hypothesis, for each s_k , either $v_i(s_k) = \mathbf{t}$ or $v_i(s_k) = \mathbf{f}$. Hence either $\varphi_{s_j}^{v_i} \equiv \top$ or $\varphi_{s_j}^{v_i} \equiv \perp$ and, consequently, either $v_{i+1}(s_j) = \mathbf{t}$ or $v_{i+1}(s_j) = \mathbf{f}$.

Since m is the maximum level of any argument in D , we now get that $v_m(s)$ is either true or false for all $s \in S$, i.e. it is a two-valued interpretation. Moreover, it holds that $v_m = \Gamma_D(v_m)$, i.e. v_m is a fixed point.

To show that v_m is the least fixed point of Γ_D , assume, towards a contradiction, that there exists an interpretation $v <_i v_m$ such that $v = \Gamma_D(v)$. Then there exists an argument s such that either $v_m(s) = \mathbf{t}$ or $v_m(s) = \mathbf{f}$, but $v(s) = \mathbf{u}$. Assume s has level i . Since D is an acyclic ADF all arguments s_k that occur in φ_s have a level less than i . Therefore, there exists at least an argument s_j of level $j < i$ in φ_s such that $v(s_j) = \mathbf{u}$. By iterating this method after at most $i - 1$ times we reach an argument s_1 of level 1 for which $v(s_1) = \mathbf{u}$. This is a contradiction, since at level 1 all arguments are initial, it must be the case that either $\varphi_{s_1} = \top$ or $\varphi_{s_1} = \perp$, and therefore $\Gamma_D(v)(s_1) \neq \mathbf{u}$. Thus, interpretation v_m is the least fixed point of Γ_D . \square

Theorem 2. *In every acyclic ADF D the sets of grounded interpretations, complete interpretations, preferred interpretations, two-valued models, and stable models coincide.*

Proof. First, the grounded interpretation v of D is also complete in D . Moreover, Proposition 1 implies that v is a two-valued model of D . Since $w = w^{\mathbf{t}} = v^{\mathbf{t}}$ in which $w = \text{grad}(D^v)$, v is a stable model. It remains to show that there is no further complete interpretation v' of D . Since v is two-valued it must hold that $v \not<_i v'$. However, since v is grounded and therefore the least complete interpretation, such a v' cannot exist. Therefore, v is the unique complete interpretation of D which is grounded, stable, two-valued, and preferred. \square

An immediate consequence of Theorem 2 is that any acyclic ADF D possesses a non-trivial preferred interpretation, which is also a complete interpretation, grounded interpretation, stable interpretation, and a model. We conclude by noting that, on the other hand, if all semantics of an ADF coincide, there is no guarantee that the ADF in question is acyclic. This is shown via Example 3.

Example 3. Consider the ADF $D = (\{a, b, c\}, \{\varphi_a = \top, \varphi_b = \neg a \wedge \neg c, \varphi_c = \neg b\})$. This ADF possesses the unique complete interpretation $v = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}$, which is also preferred, stable, grounded, as well as a model of D . That is, all semantics of D coincide; however, D is not acyclic.

The ADF of Example 3 in fact represents an AF. Therefore, there is also no guarantee that an AF is acyclic, whenever all semantics yield the same extensions.

3.2. Symmetric ADFs

We turn now to our study of symmetric ADFs. We consider the properties of coherence (stable and preferred semantics coincide) and relatively-groundedness (grounded extension is the intersection of all preferred extensions) which have been shown to hold for symmetric AFs [10]. Since both the two-valued and stable model semantics for ADFs are proper generalisations of the stable semantics for AFs [4], we consider further forms of coherence (weak coherence and semi-coherence; defined in Definition 12) that are possible in the realm of ADFs.

We will show that, contrary to symmetric AFs, symmetric ADFs do not satisfy any of the forms of coherence for ADFs we define, nor are they relatively-grounded (Theorem 3). We then define a further restricted form of symmetric ADFs, acyclic support symmetric ADFs, or ASSADFs for short (Definition 14), which we show do satisfy a weak form of coherence (each two-valued model is a stable model) (Theorem 5). Nevertheless, we conclude (Theorem 6) by showing that in ASSADFs it is still not the case that every preferred interpretation is a two-valued-model (semi-coherence). We also show that ASSADFs are not relatively-grounded (again, Theorem 6).

We start by giving the definition of symmetric ADFs.

Definition 11. An ADF $D = (S, L, C)$ is *symmetric* if L is irreflexive and symmetric and L does not contain any redundant links.

The reason why we have to exclude redundant links is that otherwise we are able to add arbitrary links without changing the semantics of the ADF at hand: informally speaking, given an ADF $D = (S, L, C)$, take any link $(a, b) \in L$ such that $(b, a) \notin L$ and do the following: add (b, a) to L and change the acceptance condition φ_a to $\varphi_a \wedge (\neg b \vee b)$. From the definition of the semantics, it follows that such a modification cannot change the set of σ -interpretations of ADF. Now, applying this modification exhaustively turns L into a symmetric relation. The added links are clearly redundant since the newly introduced parent has no semantic effect on the altered acceptance condition.

Next we provide several notions of coherence which are possible for ADFs.

Definition 12. An ADF D is called

- *coherent* if each preferred interpretation of D is a stable model of D ;
- *weakly coherent* if each two-valued model of D is a stable model of D ;
- *semi-coherent* if each preferred interpretation of D is a two-valued model of D .




Fig. 3. A symmetric ADF which is neither semi-coherent, weakly coherent nor relatively grounded.

We now turn to define the notion of relatively-groundedness for ADFs.

Definition 13. An ADF D is called *relatively grounded* if $\text{grd}(D) = \prod_i \text{prf}(D)$.

In what follows, we occasionally say that a class \mathcal{C} of ADFs is coherent (resp. semi-coherent, weakly coherent, relatively grounded) if each of its elements satisfies the respective property.

It turns out that neither of the properties analogous to those holding for symmetric AFs hold for symmetric ADFs.

Theorem 3. *The class of symmetric ADFs is neither semi-coherent, nor weakly coherent, nor relatively grounded.*

Proof. Let D be the symmetric ADF depicted in Fig. 3. It holds that $\text{prf}(D) = \{v_1, v_2\}$ with $v_1 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}, d \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$ and $v_2 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{f}, d \mapsto \mathbf{u}, e \mapsto \mathbf{u}\}$. Since v_1 is a two-valued model of D which is not stable (since $D^v = D$ and $\text{grd}(D) = \{v_{\mathbf{u}}\}$), D is not weakly coherent. Also, D is not semi-coherent since v_2 is not two-valued. In addition, $\prod_i \text{prf}(D) = v_2 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}, d \mapsto \mathbf{u}, e \mapsto \mathbf{u}\}$, but $\text{grd}(D) = \{v_{\mathbf{u}}\}$. Therefore, D is not relatively grounded. \square

Note that the ADF D used as counter-example in the proof of Theorem 3 (Fig. 3), is actually a symmetric BADF since D is a symmetric ADF in which there are no dependent links and all links are either attacking or supporting. This raises the question whether there is a particular subclass of symmetric BADF which fulfills the properties considered in Theorem 3. One natural candidate is that of acyclic support symmetric ADFs, which we present in Definition 14. For the latter we first need to define the notion of a support cycle.

Definition 14. Given an ADF $D = (S, L, C)$, let L^+ be the set of all supporting links in D . An ADF D is an *acyclic support symmetric ADF* (ASSADF for short) if it is symmetric, bipolar and (S, L^+) is acyclic.

The method of determining whether a given ADF is an ASSADF is clarified in Example 4.

Example 4. Let $D = (S, L, C)$ be the symmetric BADF depicted in Fig. 3. Since (S, L^+) as shown in Fig. 4 contains a cycle, namely the sequence $[d, c, d]$, D is not an ASSADF.

We now show that ASSADF are weakly coherent, using the following technical lemma.

Lemma 4. *Let D be an ADF, v a two-valued model of D , and $s \in S$ an argument such that all parents of s are attackers. If φ_s^v is irrefutable then $\varphi_s[s_i/\perp : v(s_i) = \mathbf{f}]$ is irrefutable.*



Fig. 4. The support-links of the ADF D from Theorem 3 (Figure 3).

Fig. 5. An ASSADF without supporting links that is not semi-coherent.

Theorem 5. *Every acyclic support symmetric ADF (ASSADF) is weakly coherent.*

Proof. Let $D = (S, L, C)$ be an acyclic support symmetric ADF. We have to show that each two-valued model of D is also a stable model of D . Let hence v be a two-valued model of D , $D^v = (S^v, L^v, C^v)$ be the *stb*-reduct of D , w be the unique grounded interpretation of D^v , and φ'_s the acceptance condition of s in D^v , i.e. $\varphi'_s = \varphi_s[s_i/\perp : v(s_i) = \mathbf{f}]$. We show that $v^{\mathbf{t}} = w^{\mathbf{t}}$. Suppose to the contrary that there exists an argument s , such that $v(s) = \mathbf{t}$ and $w(s) \neq \mathbf{t}$. That is, φ'_s is not irrefutable. This means, by Lemma 4, that φ'_s contains an argument s_1 supporting s such that $v(s_1) = \mathbf{t}$, otherwise φ'_s cannot be irrefutable. Thus, φ'_s also contains s_1 . Since supports are acyclic in ASSADFs, for the same reasons $\varphi'_{s_1} = \varphi_{s_1}[s_i/\perp : v(s_i) = \mathbf{f}]$ contains an argument s_2 which is different from s and s_1 and which supports s_1 . Thus there exists an infinite sequence of arguments s_1, s_2, \dots such that s_{i+1} supports s_i . This is a contradiction to D being an ASSADF. \square

We conclude this section by showing, on the other hand, that there are ASSADFs which are neither semi-coherent nor relatively grounded.

Theorem 6. *The class of ASSADFs is neither semi-coherent nor relatively grounded.*

Proof. Consider the ASSADF D depicted in Fig. 5. D has 4 preferred interpretations, namely $v_1 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}, d \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$, $v_2 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, d \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$, $v_3 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{f}, d \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$, and $v_4 = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}, d \mapsto \mathbf{f}, e \mapsto \mathbf{t}\}$. As every two-valued interpretation of D (that is v_1, v_2 and v_3) is also a stable model, D is weakly coherent, confirming Theorem 5. However, v_4 is a preferred interpretation which is not a two-valued model. Hence, D is not semi-coherent.

We show now that ASSADFs are not relatively grounded in general. Consider the ASSADF $D = (S, L, C)$, depicted in Fig. 6. Here $D = (\{a, b, c\}, \{\varphi_a : \neg b \wedge \neg c, \varphi_b : \neg a \wedge \neg c, \text{ and } \varphi_c : a \vee \neg b\})$. D has the preferred interpretations $v_1 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}$ and $v_2 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}\}$. We obtain $v_1 \sqcap_i v_2 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}\}$. However, the grounded interpretation of D is the trivial interpretation v_u . That is, D is not relatively grounded. \square




Fig. 6. An ASSADF which is not relatively grounded.

3.3. Implications for SETAFs

The ASSADF used in the proof of Theorem 6 to show that ASSADFs are not semi-coherent does not have any supporting links. That is, even ASSADFs without supporting links are not semi-coherent. This leads us, in this section, to consider whether ASSADFs having only attacking links, which we call support free symmetric ADFs or SFSADFs for short (Definition 15), satisfy the other properties considered in Section 3.2: being weakly coherent and relatively grounded. We show that SFSADFs are weakly coherent, but neither semi-coherent nor relatively grounded in Theorem 8.

Moreover, we derive from Theorem 8 that symmetric SETAFs are neither coherent nor relatively grounded (Theorem 17). The reason is that the SFSADFs we have used in the proof of Theorem 8 correspond to SETAFs. More concretely, the SETAFs in question correspond to a specific class of SFSADFs: those in which the acceptance condition of none of the arguments is unsatisfiable. On the way of proving Theorem 17 we show that, in fact, such SFSADFs exactly correspond to symmetric SETAFs (Theorem 9, Corollary 10, and Theorem 12; Lemmas 13 and 14). Thus, we obtain as a consequence of our investigations of semantic properties in the general settings of ADFs, results that complement those of [27] for SETAFs, where the authors show that the standard semantics are indistinguishable on acyclic SETAFs (a result that is confirmed by our study in Section 3.1).

We start by defining SFSADFs:

Definition 15. Given an ADF $D = (S, L, C)$, let L^- be the set of all attacking links in D . A bipolar ADF $D = (S, L, C)$ is a *support free symmetric ADF* (SFSADF for short) if it is symmetric and does not have any supporting links, that is, $L = L^-$.

Note that since SFSADFs are BADFs by Definition 15, this means that SFSADFs do not have dependent links. Also, SFSADFs are symmetric by the same definition, which means that they do not have redundant links. Further, since SFSADFs do not have supporting links, they also do not have a support cycle. Thus, the class of SFSADFs is indeed a strict subclass of ASSADFs.

We next show that also SFSADFs, while being weakly coherent, are neither semi-coherent nor relatively grounded. Before doing so, we report a simple observation concerning the grounded interpretation.

Lemma 7. *Let D be an SFSADF with no isolated argument. The unique grounded interpretation of D is the trivial interpretation, $v_{\mathbf{u}}$.*

Proof. We show that for any SFSADF $D = (S, L, C)$ with no isolated argument, $\Gamma_D(v_{\mathbf{u}}) = v_{\mathbf{u}}$. Let s be an argument. Let v_1 be an interpretation in which all parents of s are assigned to \mathbf{t} and let v_2 be an interpretation in which all $\text{par}(s)$ are assigned to \mathbf{f} . Since D is an SFSADF, the former interpretation shows that $\varphi_s^{v_{\mathbf{u}}}$ is not irrefutable and the latter interpretation says that $\varphi_s^{v_{\mathbf{u}}}$ is not unsatisfiable. Therefore, for each argument s , $\Gamma_D(v_{\mathbf{u}})(s) = \mathbf{u}$. \square




Fig. 7. A SFSADF that is not relatively grounded.

Theorem 8. *The following are the case for SFSADFs:*

- every SFSADF is weakly coherent,
- the class of SFSADFs is not semi-coherent,
- the class of SFSADFs is not relatively grounded.

Proof. By Theorem 5, every ASSADF is weakly coherent, and since each SFSADF is an ASSADF, also SFSADFs are weakly coherent.

The ASSADF D , depicted in Fig. 5 to show that the class of ASSADFs is not semi-coherent, does not have any supporting links. That is, D is a SFSADF. Thus, the class of SFSADFs is not semi-coherent either.

It remains to show that the class of SFSADFs is not relatively grounded. Let D be the ADF depicted in Fig. 7. The unique preferred interpretation of D is $v = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{f}, d \mapsto \mathbf{t}\}$. However, since D does not possess an isolated argument, Lemma 7 shows that the grounded interpretation of D is the trivial interpretation. That is, the meet of the preferred interpretations of D is not equal to the grounded interpretation of D . Hence, D is not relatively grounded. \square

It is relatively easy to see that the ADFs used to show that SFSADFs are neither semi-coherent nor relatively grounded in the proof of Theorem 8 (ADFs from Figs 5 and 7) correspond to SETAFs (see Definitions 3 and 9). In fact, we proceed to show now that symmetric SETAFs are captured exactly by a subclass of SFSADFs: those in which the acceptance condition of none of the arguments is unsatisfiable. As already hinted at, apart from showing the link between SETAFs and SFSADFs this will allow us to also formally translate the content of Theorem 8 to the context of SETAFs.

We start by defining symmetric SETAFs.

Definition 16. A SETAF $F = (A, R)$, in which $R \subseteq (2^{\{\emptyset\}} \times A)$, is a *symmetric SETAF* if the following properties hold:

- for all $(S, t) \in R$ and for all $s \in S$, there exists $(T, s) \in R$ such that $t \in T$,
- for each argument s and for each $(S, s) \in R$, the set S does not include s ,
- for each $(S, s) \in R$ there is no $(S', s) \in R$ with $S' \subset S$.

In Definition 16, the first item indicates that in the associated graph of the symmetric SETAFs all links are symmetric. The second item further means that there are also no reflexive links. Finally, the third item excludes redundant links.

From Definition 9 it follows that each SETAF can be represented as an ADF. Thus, also symmetric SETAFs can be encoded as ADFs. We now show that in fact symmetric SETAFs correspond to SFSADF.

Theorem 9. *The ADF associated to a given symmetric SETAF is a SFSADF.*

Proof. Let $F = (A, R)$ be a symmetric SETAF. We show that the ADF $D_F = (S, L, C)$ associated to F is a SFSADF. By Definition 9, D_F does not contain any supporting link. It remains to show that D_F is a symmetric ADF. By Definition 16, it is clear that L does not have any redundant links. We hence show that L is symmetric and irreflexive. Towards a contradiction, thus assume that either L is not symmetric or not irreflexive.

- Assume that L is not symmetric. This means that there is an argument s which is a parent of t but not visa versa. That is, s appears in the acceptance condition of t but t does not appear in the acceptance condition of s . Since s is a parent of t , by Definition 9 there is $(S, t) \in R$ such that $s \in S$. Since F is a symmetric SETAF, there is a $(T, s) \in R$ such that $t \in T$. Then, again via Definition 9, the argument t appears in the acceptance condition of s . Thus, t is a parent of s . This shows that the assumption that L is not symmetric is false.
- Assume now that L is not irreflexive. Therefore, there is an argument s which is contained in $par(s)$. By Definition 9 there is $(S, s) \in R$ such that $s \in S$ which is in contradiction to F being symmetric, cf. Definition 16.

Thus, if F is a symmetric SETAF, then the associated ADF D_F is a SFSADF. \square

Let $F = (A, R)$ be a SETAF and let a be an argument on which there is no attack in F , that is, there is no $(B, a) \in R$. By Definition 9, in the ADF corresponding to the SETAF F the acceptance condition of a has the form $\varphi_a = \bigwedge_{(B,a) \in R} \bigvee_{a' \in B} \neg a' \equiv \top$. On the other hand, if there exists $(B, a) \in R$, it holds that $\varphi_a = \bigwedge_{(B,a) \in R} \bigvee_{a' \in B} \neg a' \neq \perp$. These facts together with Theorem 9 lead to the following corollary.

Corollary 10. *The SFSADF associated to a given symmetric SETAF does not contain any argument with an unsatisfiable acceptance condition.*

Next, we detail how the special group of SFSADFs that do not have any argument with an unsatisfiable acceptance condition can be represented as SETAFs. For this we make use of a fact from [36], namely that ADFs for which the acceptance condition of each argument is either tautological or in CNF having only negative literals can be represented as SETAFs. Note that in symmetric ADFs each initial argument needs to be isolated and thus might have acceptance condition \top or \perp ; the latter is problematic in representing SFSADFs as symmetric AFs and thus needs special treatment.

Lemma 11. *Let $D = (S, L, C)$ be a SFSADF and let $s \in S$ be an argument that is not isolated. Then the acceptance condition of the argument s can be written in conjunctive normal form and having only negative literals.*

Proof. Since the acceptance condition of each argument in an ADF is indicated by a propositional formula, it can be transformed to CNF. It remains to show that each of the resulting formulas in CNF consists of only negative literals. Toward a contradiction assume that φ_s is the acceptance condition of an argument s , that is not isolated, in conjunctive normal form that contains positive literals. Without

loss of generality assume that t is the only argument that appears in φ_s as a positive literal. Let $\{c_t^i\}_{1 \leq i \leq n}$ be the set of all clauses c_t^i in which t occurs ($n \geq 1$). Also, let v be a two-valued interpretation which assigns the truth value \mathbf{f} to t ; all other arguments in each c_t^i is assigned to \mathbf{t} . Also v assigns \mathbf{f} to all other arguments of $\text{par}(s)$, which means that $v(\varphi_s) = \mathbf{f}$. However, $v|_{\mathbf{t}}(\varphi_s) = \mathbf{t}$. Thus, by the definition of attacking links, (t, s) is not an attacking link in D . This is a contradiction with D being a SFSADF. \square

We now provide the construction associating a SETAF to every SFSADF for which no argument has an unsatisfiable acceptance condition.

Definition 17. Let $D = (S, L, C)$ be a SFSADF in which there is no argument having an unsatisfiable acceptance condition. D can be written as a SETAF $F_D = (A, R)$ such that $A = S$ and R is as follows. Let φ_a be in CNF having only negative literals, let c be a clause of φ_a and let R_c be the set of all arguments in the clause c . Then, (R_c, a) represents a joint attack to a . The set $R_a = \{(R_c, a) | c \text{ is a clause in } \varphi_a\}$ is the set of all joint attacks to an argument a . Let $R = \bigcup_{a \in S} R_a$, be the set of all joint attacks in F_D . We call F_D the SETAF associated to D .

Analogously to Theorem 9, we show in Theorem 12 that SFSADFs in which none of the acceptance conditions of the arguments is unsatisfiable can be mapped to symmetric SETAFs.

Theorem 12. Let $D = (S, L, C)$ be a SFSADF in which the acceptance condition of none of the arguments is unsatisfiable. The SETAF F_D associated to D is a symmetric SETAF.

Proof. Assume that D is a SFSADF in which there is no argument with an unsatisfiable acceptance condition. Thus, via Definition 17, D can be written as a SETAF $F_D = (A, R)$. It remains to show that F_D is a symmetric SETAF. Towards a contradiction, assume that F_D is not symmetric. Then, there are two possibilities to consider:

- There exists $(S, t) \in R$ and there exists $s \in S$ such that for all $(T, s) \in R$, T does not include t . Therefore, the acceptance condition of t in D includes s , however, the acceptance condition of s in D does not include t . That is, L in D is not symmetric.
- There exists $(S, t) \in R$ such that $t \in S$. That is, t appears in the acceptance condition of t in D . That is, L in D contains a reflexive link.

Both possibilities are in contradiction with the assumption that D is a SFSADF. Therefore, the assumption that F_D is not symmetric is not true. Then, the SETAF associated to a SFSADF is a symmetric SETAF. \square

The SFSADF $D = (\{a, b, c, d\}, \{\varphi_a = \neg b \wedge \neg c, \varphi_b = \neg a \wedge \neg c \wedge \neg d, \varphi_c = \neg a \wedge \neg b \wedge \neg d, \varphi_d = \neg b \vee \neg c\})$, depicted in Fig. 7, corresponds to the symmetric SETAF $F_D = (\{a, b, c, d\}, R)$ with $R = (\{\{a\}, b\}, (\{b\}, a), (\{a\}, c), (\{c\}, a), (\{b\}, c), (\{c\}, b), (\{b, c\}, d), (\{d\}, b), (\{d\}, c))$ depicted in Fig. 8. As can be seen in the figure, in this SETAF there is a joint attack from b and c to d . This joint attack is symmetric in the sense of Definition 16, because of $(\{d\}, b) \in R$ and $(\{d\}, c) \in R$.

Now, in order to be able to relate SFSADFs and symmetric SETAFs on the semantic level, we first make precise the relation between extension-based semantics of AFs and interpretation-based semantics of ADFs. To do so, we start by introducing some terminology. Given a formalism F , the set of all extensions of F are denoted by \mathcal{E} and the set of all possible interpretations of F are denoted by \mathcal{V} . The function Ext2Int_F , in Definition 18, is a modification of the function (associating labellings to extensions) given in Definition 5.1. of [21].




Fig. 8. A symmetric SETAF that is not relatively grounded.

Definition 18. Let $F = (A, R)$ be a SETAF, and let e be an extension of F ($e \in \mathcal{E}$). The truth value assigned to each argument $a \in A$ by the three-valued interpretation v_e associated to e is given by the function $Ext2Int_F : \mathcal{E} \rightarrow \mathcal{V}$ as follows.

$$Ext2Int_F(e)(a) = \begin{cases} \mathbf{t} & a \in e, \\ \mathbf{f} & \exists B \subseteq 2^A \text{ such that } (B, a) \in R \text{ and } \forall b \in B, b \in e, \\ \mathbf{u} & \text{otherwise.} \end{cases}$$

An ADF interpretation, on the other hand, can be represented as an extension via the following mapping.

Definition 19. Let $D = (S, L, C)$ be an ADF and v an interpretation of D , that is, $v \in \mathcal{V}$. The associated extension e_v of v is obtained via application of the function $Int2Ext_D : \mathcal{V} \rightarrow \mathcal{E}$ on v , as follows:

$$Int2Ext_D(v) = \{s \in S \mid s \mapsto \mathbf{t} \in v\}$$

The two subsequent lemmas are adopted from [21].

Lemma 13. Let F be a SETAF, and $\sigma \in \{prf, stb, mod, com, grd\}$. Moreover, let D_F be the ADF associated to F . Then, $\sigma(D_F) = \{Ext2Int_F(e) \mid e \in \sigma(F)\}$.

Lemma 14. Let D be a SFSADF in which the acceptance condition of none of the arguments is unsatisfiable, $\sigma \in \{adm, prf, stb, mod, com, grd\}$ and F_D the SETAF associated to D . Then, $\sigma(F_D) = \{Int2Ext_D(v) \mid v \in \sigma(D)\}$.

Note that Lemma 13 does not mention admissible semantics, while Lemma 14 does. The reason is that for a given SETAF F the associated three-valued interpretations obtained via $Ext2Int_F$ usually do not cover all admissible interpretations of the ADF D_F . This is illustrated next.

Example 5. Let $F = (\{a, b, c\}, \{\{a, b, c\}\})$ be a SETAF. By Definition 9, the associated ADF to F is $D_F = (\{a, b, c\}, \{\varphi_a = \top, \varphi_b = \top, \varphi_c = \neg a \vee \neg b\})$. It is clear that $e = \{a, b\} \in adm(F)$. Applying $Ext2Int_F(e)$ of Definition 18 to e leads to the three-valued interpretation $v_e = \{a \mapsto t, b \mapsto t, c \mapsto f\}$.

However, $\{a \mapsto t, b \mapsto t, c \mapsto u\}$ is also an admissible interpretation of D_F which is not obtained from any admissible extension e of F via $Ext2Int_F(e)$.

One can overcome this problem by mapping each admissible set e of a SETAF $F = (A, R)$ to several interpretations in a way that $Ext2Int_F(e)(a)$ yields either \mathbf{u} or \mathbf{f} in case there exists $(B, a) \in R$ such that $B \subseteq e$, and for all (B', c) with $a \in B'$ and $c \in e$ there exists $(b \neq a) \in B'$ such that $Ext2Int_F(e)(b) = \mathbf{f}$. That is, $Ext2Int_F(e)(a)$ can be either \mathbf{f} or \mathbf{u} , if a is attacked by some arguments of e but the truth value of a does not play any role for the truth value of elements of e . However, since the forthcoming results do not involve admissible semantics, we leave a more formal investigation on this issue as topic for future work.

The next lemma rephrases Lemma 7 in terms of SETAFs.

Lemma 15. *Let F be a symmetric SETAF with no isolated argument. The unique grounded extension of F is the empty set.*

Proof. Let $F = (A, R)$ be a symmetric SETAF with no isolated argument, and let $D_F = (S, L, C)$ be the associated ADF of F . Toward a contradiction assume that the unique grounded extension e (in F) is not the empty set, that is, there exists argument a such that $a \in e$. First, we show that D_F does not contain any isolated argument. To this end, let b be an argument. Since F does not contain any isolated argument, there exists $B \subseteq A$ such that $(B, b) \in R$. The associated acceptance condition of b in D_F , namely $\varphi_b = \bigwedge_{(B,b) \in R} \bigvee_{b' \in B} \neg b'$ shows that $par(b) \neq \{\}$ in D_F . Thus, the associated D_F does not contain any isolated argument, as well. By Lemma 13 and Definition 18, a is assigned to \mathbf{t} in $grd(D_F)$. Further by Theorem 9, D_F is an SFSADF. This is a contradiction by Lemma 7. Therefore, the unique grounded extension of F is the empty set. \square

The following corollary is a direct result of Lemma 15.

Corollary 16. *A symmetric SETAF F with no isolated argument is relatively grounded if and only if the intersection of all preferred extensions of F is the empty set.*

We are now in position to derive that symmetric SETAFs are neither coherent nor relatively grounded from our proof of Theorem 8.

Theorem 17. *The class of symmetric SETAFs is neither coherent nor relatively grounded.*

Proof. The ADFs which are used in the proof of Theorem 8, to show that the class of SFSADFs is neither semi-coherent (and thus not coherent) nor relatively grounded, do not consist of any argument with an unsatisfiable acceptance condition. Then, by Theorem 12, the associated AFs to those SFSADFs are symmetric SETAFs.

Now, let D be such an SFSADF that does not satisfy coherence. We show that the SETAF F_D associated to D cannot be coherent either. Let w be a preferred interpretation of D that is not a stable model of D . By Lemma 14, $prf(F_D) = \{Int2Ext_D(v) \mid v \in prf(D)\}$ and $stb(F_D) = \{Int2Ext_D(v) \mid v \in stb(D)\}$. Towards a contradiction, suppose $prf(F_D) = stb(F_D)$. It follows that there is an interpretation $u \in prf(D)$ with $u^t = w^t$, such that $u \in stb(D)$, and hence $u \neq w$. Hence, either $u^f \subset w^f$ or $w^f \subset u^f$. In the first case $u <_i w$ and hence u cannot be preferred; in the second case $w <_i u$ and hence w cannot be preferred. In both cases we have a contradiction.

For relatively groundedness, we already have provided a symmetric SETAF that violates this property in Fig. 8. For that SETAF it can be checked that its complete sets are \emptyset and $\{a, d\}$. Corollary 16 immediately implies that this SETAF cannot be relatively grounded. \square

3.4. The role of odd-length cycles

In [15] it is proven that if an AF is not coherent then it contains a cycle of odd length. This means, on the other hand, that if an AF does not contain any odd-length cycle it is coherent. It is easy to show that this property does not generalise to ADFs because of the possibility of support links. Indeed, consider the ADF $D = (\{a, b, c, d\}, \{\varphi_a = d, \varphi_b = a, \varphi_c = b, \varphi_d = c\})$. The interpretation $v = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}, d \mapsto \mathbf{t}\}$ is a two-valued (and, hence, preferred) model of D , however it is not a stable model. That is, D is not weakly coherent and, therefore, also not coherent.

In this section we study whether the property of having only odd-length-cycles implying coherence carries over to any subclasses of ADFs we introduced in our work so far. In Section 3.2 we had shown that ASSADFs are weakly coherent but not semi-coherent. We here show that also ASSADFs not containing any odd-length-cycle are not semi-coherent and, thus, that such ASSADFs are not coherent (Theorem 18). On the other hand, we are able to show that SFSADFs not containing any odd-length-cycle are coherent (Theorem 19). In fact we prove a more general result: bipolar ADFs without supporting links, which we dub SFADFs (Definition 20; note the difference with SFSADFs which have an ‘‘S’’ after the first ‘‘F’’), that also do not have any odd-length-cycle are coherent (Corollary 20). Moreover, given that SETAFs correspond to a special class of SFSADFs (see Section 3.3), the result also applies to SETAFs.

Theorem 18. *The subclass of ASSADFs not containing any odd-length cycle is not coherent.*

Proof. Consider the ADF $D = (\{a, b\}, \{\varphi_a = b, \varphi_b = \neg a\})$. D is an ASSADF in which there is no odd cycle. The unique preferred interpretation of D is $v = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}$, however, v is not a two-valued model. Then, D is not semi-coherent and D is not coherent either. \square

Contrarily, we show next that the subclass of SFSADFs in which each ADF does not contain any odd cycle is coherent.

Theorem 19. *The class of SFSADFs that do not contain any odd-length cycle is coherent.*

Proof. Let $D = (S, L, C)$ be a SFSADF that does not contain any odd cycle. Since we showed in Theorem 8 that the class of SFSADFs is weakly coherent, D is weakly coherent, as well. To complete the proof of the theorem it is enough to show that D is semi-coherent, i.e. $\text{prf}(D) = \text{mod}(D)$. Since each two-valued model is a preferred interpretation of ADFs, it is trivial that $\text{prf}(D) \supseteq \text{mod}(D)$. Thus, it is enough to show that $\text{prf}(D) \subseteq \text{mod}(D)$. Toward a contradiction, assume that there exist a preferred interpretation v of D that is not a two-valued model. That is, there must exist an argument a such that $v(a) = \mathbf{u}$.

Let S be the set of all arguments which are assigned the truth value \mathbf{u} by v . Since D is a SFSADF that does not contain any odd cycle, the arguments of S have to be in even cycles in the associated graph. Without loss of generality, assume that the associated graph of D contains only one such even cycle. Then one can construct a bipartite graph of nodes of this cycle with partitions S_1 and S_2 . Assign all

arguments in S_1 to \mathbf{t} , and all arguments of S_2 to \mathbf{f} . Construct the interpretation v' as follows.

$$v'(a) = \begin{cases} v(a) & \text{if } v(a) = \mathbf{t}/\mathbf{f}, \\ \mathbf{t} & \text{if } a \in S_1, \\ \mathbf{f} & \text{if } a \in S_2. \end{cases}$$

It is clear that $v \leq_i v'$. We now show that v' is a two-valued model. First, there is no argument in v' assigned to \mathbf{u} . To show that v' is a two-valued model it remains to show that $\Gamma_D(v') = v'$. Assume that a is assigned to \mathbf{t} in v' , we show that $\Gamma_D(v')(a) = \mathbf{t}$. (The method for proving the case that a is assigned to \mathbf{f} is analogous.) If $a \mapsto \mathbf{t}$ in v' either $v(a) = \mathbf{t}$ or $a \in S_1$.

- If $v(a) = \mathbf{t}$, since v is a preferred interpretation, $\Gamma_D(v)(a) = \mathbf{t}$. In addition, the characteristic operator is a monotonic operator, that is, $\Gamma_D(v')(a) = \mathbf{t}$.
- If $a \in S_1$, then $\text{par}(a) \neq \{\}$. Let φ_a be in CNF, this is possible by Lemma 11.
 - * If $\Gamma_D(v')(a) = \mathbf{f}$, then there exists a clause c in φ_a all arguments of which are assigned to \mathbf{t} in v' . By the construction of S_1 , $\text{par}(a) \not\subseteq S_1$. Therefore, all arguments of c are assigned to \mathbf{t} in v (note that all arguments in S_2 are assigned to \mathbf{f} by the definition of v'). That is, $\Gamma_D(v)(a) = \mathbf{f}$, which is a contradiction with the assumption that $v(a) = \mathbf{u}$. Thus, $\Gamma_D(v')(a) \neq \mathbf{f}$.
 - * If $\Gamma_D(v')(a) = \mathbf{u}$, then there exists a parent of a the truth value of which is not indicated in v' . This is also a contradiction, since all parents of a are either in S_2 that are assigned to \mathbf{f} in v' or assigned to \mathbf{t}/\mathbf{f} by v . Therefore, $\Gamma_D(v')(a) \neq \mathbf{u}$.

Thus, if $a \in S_1$, then $\Gamma_D(v')(a) = \mathbf{t}$.

Therefore, v' is a two-valued model of D and hence a preferred interpretation of D . Moreover $v \leq_i v'$, which is a contradiction to the assumption that v is a preferred interpretation. Therefore, each preferred interpretation of a SFSADF that does not contain any odd cycle is a two-valued model. Thus, if a SF-SADF does not contain any odd cycle, then it is coherent. \square

As it turns out, the proof of Theorem 19 is independent from the notion of symmetry. Hence, we obtain as a final observation (Corollary 20) in this section that the general class of support-free ADFs which we define next is coherent.

Definition 20. Given an ADF $D = (S, L, C)$, let L^- be the set of all attacking links in D . A bipolar ADF $D = (S, L, C)$ is called a support-free ADF (SFADF) if it does not have any supporting links, that is, $L = L^-$.

Corollary 20. *SFADFs and SETAFs without odd-length cycles are coherent.*

4. Expressiveness of ADF subclasses

Following the work of [16] in this section we consider the expressiveness, i.e. the set of possible outcomes which can be achieved under the different semantics, of ASSADFs and SFSADFs. We thus complement here results that have been obtained for general [29,33] and bipolar ADFs [25] and also compare the ADF subclasses we introduced in this work to abstract argumentation frameworks in terms of their expressivity.

Formally, the study of expressivity of a formalism w.r.t. a semantics can be done by considering the outcomes that can be realised by the formalism under the semantics of interest.

Definition 21. Let \mathcal{F} be a formalism (e.g. AFs or (subclasses of) ADFs), i.e. the set of structures available in \mathcal{F} (e.g. all possible AFs and ADFs) and σ a semantics for \mathcal{F} . Moreover, let \mathbb{V} be an interpretation-set or extension-set. \mathbb{V} is said to be σ -realizable in \mathcal{F} , if there exists an element kb (“knowledge base”) of \mathcal{F} such that $\sigma(\text{kb}) = \mathbb{V}$.

The signature of a formalism w.r.t. a semantics is then the set of possible outcomes that can be realised by the formalism under the semantics, this is encoded in Definition 22.

Definition 22. The *signature* $\Sigma_{\mathcal{F}}^{\sigma}$ of a formalism \mathcal{F} w.r.t. a semantics σ is defined as:

$$\Sigma_{\mathcal{F}}^{\sigma} = \{\sigma(\text{kb}) \mid \text{kb} \in \mathcal{F}\}.$$

Formalisms can now be compared for expressivity by considering their signatures. Specifically, given two formalisms $\mathcal{F}_1, \mathcal{F}_2$ as well as a semantics σ , we say that \mathcal{F}_1 is *strictly more expressive* than \mathcal{F}_2 for σ , whenever $\Sigma_{\mathcal{F}_2}^{\sigma} \subsetneq \Sigma_{\mathcal{F}_1}^{\sigma}$. \mathcal{F}_1 and \mathcal{F}_2 are, on the other hand, incomparable under a semantics σ if neither $\Sigma_{\mathcal{F}_1}^{\sigma} \subseteq \Sigma_{\mathcal{F}_2}^{\sigma}$ nor $\Sigma_{\mathcal{F}_2}^{\sigma} \subseteq \Sigma_{\mathcal{F}_1}^{\sigma}$. This is denoted as $\Sigma_{\mathcal{F}_1}^{\sigma} \approx \Sigma_{\mathcal{F}_2}^{\sigma}$.

In what follows we concentrate on studying ASSADFs and SFSADFs from the perspective of realizability. We compare these novel subclasses of ADFs to that of AFs, BADFs, and general ADFs. We build on studies comparing the expressivity of AFs, BADFs and (general) ADFs reported on in [25,32].

We begin by showing that BADFs are strictly more expressive than ASSADFs for the admissible, preferred, complete, and model semantics.

Theorem 21. For $\sigma \in \{\text{adm}, \text{prf}, \text{com}, \text{mod}\}$ it holds that $\Sigma_{\text{ASSADF}}^{\sigma} \subsetneq \Sigma_{\text{BADF}}^{\sigma}$.

Proof. Since every ASSADF is, by definition, a BADF, $\Sigma_{\text{ASSADF}}^{\sigma} \subseteq \Sigma_{\text{BADF}}^{\sigma}$ is clear. To show that $\Sigma_{\text{BADF}}^{\sigma}$ is a strict superset of $\Sigma_{\text{ASSADF}}^{\sigma}$ it is enough to find an interpretation-set \mathbb{V} which is σ -realizable in BADFs, but not σ -realizable in ASSADFs, for $\sigma \in \{\text{adm}, \text{prf}, \text{com}, \text{mod}\}$.

For $\sigma \in \{\text{prf}, \text{mod}\}$, let $\mathbb{V} = \{\{a \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{f}\}\}$, and for $\sigma' \in \{\text{com}, \text{adm}\}$, let $\mathbb{V}' = \{\{a \mapsto \mathbf{u}\}, \{a \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{f}\}\}$. The BADF $D = (S, L, C)$ with $S = \{a\}$ and $\varphi_a = a$ realizes \mathbb{V} under σ and \mathbb{V}' under σ' . On the other hand, it is easy to check that there is no ASSADF with one argument which realizes \mathbb{V} under σ , and respectively, \mathbb{V}' under σ' . To complete the proof toward a contradiction assume that there exists an ASSADF $D' = (S', L', C')$ such that $\sigma(D') = \mathbb{V}$ and $\sigma'(D') = \mathbb{V}'$. Since \mathbb{V} contains an assignment to only one argument, D' has to have just one argument. Since D' is an ASSADF, the argument a is an isolated argument in D' . That is, either $\varphi_a = \top$ or $\varphi_a = \perp$. Thus, it can realize neither \mathbb{V} under σ nor \mathbb{V}' under σ' . \square

Next, we show that ASSADFs are strictly more expressive than SFSADFs for the admissible, preferred, and complete semantics. In a certain sense, this complements our observation in Section 3.4 that ASSADFs and SFSADFs differ in terms of coherence on odd-cycle free frameworks.

Theorem 22. For $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$, it holds that $\Sigma_{\text{SFSADF}}^{\sigma} \subsetneq \Sigma_{\text{ASSADF}}^{\sigma}$.

Proof. Since each support-free symmetric ADF is an acyclic support symmetric ADF, it is clear that $\Sigma_{\text{SFSADF}}^\sigma \subseteq \Sigma_{\text{ASSADF}}^\sigma$. To show that $\Sigma_{\text{ASSADF}}^\sigma$ is a strict superset of $\Sigma_{\text{SFSADF}}^\sigma$, for $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$, we give an interpretation-set \mathbb{V} , which is σ -realizable in ASSADFs, but not σ -realizable in SFSADFs.

Let $\mathbb{V} = \{\{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}\}$. The ASSADF $D = (S, L, C)$ with $S = \{a, b\}$, $\varphi_a = b$, and $\varphi_b = \neg a$ realizes \mathbb{V} under $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$. However, it is easy to check that there is no SFSADF that can realize \mathbb{V} . Assume to the contrary that there exists a SFSADF $D' = (S', L', C')$ that can realize \mathbb{V} under σ , for $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$. The set of arguments of D' is $\{a, b\}$ as these are the only ones appearing in the σ interpretations of D' .

- If any of the arguments of S' is an isolated argument, then its acceptance condition is either equivalent to \top or \perp . That is, $\sigma(D') \neq \mathbb{V}$. Thus, none of the arguments could be an isolated argument in D' .
- If there is a symmetric attack relation between a and b , then the interpretation $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}\}$ is a preferred interpretation of D' and therefore it is a complete and admissible interpretation of D' . Therefore, $\sigma(D') \neq \mathbb{V}$, for $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$.

Thus, \mathbb{V} is not σ -realizable in SFSADFs under $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$. \square

The forthcoming result shows why, on the other hand, AFs and SFSADFs are incomparable, and also AFs and ASSADFs are incomparable, in terms of their expressivity for the admissible, preferred, and complete semantics.


Theorem 23. For $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$, it holds that $\Sigma_{\text{AF}}^\sigma \approx \Sigma_{\text{SFSADF}}^\sigma$ and $\Sigma_{\text{AF}}^\sigma \approx \Sigma_{\text{ASSADF}}^\sigma$.

Proof. To obtain our theorem we show that $\Sigma_{\text{AF}}^\sigma \not\subseteq \Sigma_{\text{ASSADF}}^\sigma$ and $\Sigma_{\text{SFSADF}}^\sigma \not\subseteq \Sigma_{\text{AF}}^\sigma$. Since, via Theorem 22, $\Sigma_{\text{ASSADF}}^\sigma$ is a strict superset of $\Sigma_{\text{SFSADF}}^\sigma$ for $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$, we can then conclude that $\Sigma_{\text{AF}}^\sigma \not\subseteq \Sigma_{\text{SFSADF}}^\sigma$ and $\Sigma_{\text{ASSADF}}^\sigma \not\subseteq \Sigma_{\text{AF}}^\sigma$.

• To show $\Sigma_{\text{AF}}^\sigma \not\subseteq \Sigma_{\text{ASSADF}}^\sigma$ consider $\mathbb{V} = \{\{a \mapsto \mathbf{u}\}\}$. A witness of σ -realizability in AFs is $F = (\{a\}, \{(a, a)\})$. However, there is no ASSADF to realize \mathbb{V} under σ .

• To verify that $\Sigma_{\text{SFSADF}}^\sigma \not\subseteq \Sigma_{\text{AF}}^\sigma$ for $\sigma \in \{\text{adm}, \text{prf}, \text{com}\}$ we first show that $\Sigma_{\text{SFSADF}}^{\text{prf}} \not\subseteq \Sigma_{\text{AF}}^{\text{prf}}$. Let $\mathbb{V} = \{v_1, v_2, v_3\}$ with $v_1 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}, e \mapsto \mathbf{t}\}$, $v_2 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$, and $v_3 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, e \mapsto \mathbf{f}\}$. A witness of *prf*-realizability of \mathbb{V} in SFSADFs is $D = (S, L, C)$ with $S = \{a, b, c, e\}$, $\varphi_a = \neg e \wedge (\neg b \vee \neg c)$, $\varphi_b = \neg a \vee \neg c$, $\varphi_c = \neg a \vee \neg b$, and $\varphi_e = \neg a$. However, there is no AF with \mathbb{V} as its preferred interpretations. If there is an AF F' such that $\sigma(F') = \mathbb{V}$ then the structure of v_1, v_2 and v_3 implies that there is no attack between a, b and c in F' . Thus, if there is an attack from any of a, b and c to e then $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}, e \mapsto \mathbf{f}\}$ is a preferred interpretation of F' . If there is no attack from any of a, b and c to e then $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}, e \mapsto \mathbf{t}\}$ is a preferred interpretation of F' . In both cases $\sigma(F') \neq \mathbb{V}$. For $\sigma = \text{com}$ let $\mathbb{V}' = \mathbb{V} \cup \{\{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}, e \mapsto \mathbf{u}\}\}$. It is easy to check that \mathbb{V}' is *com*-realizable by the SFSADF D defined above. If there is an AF F' that realizes \mathbb{V}' under *com* then each of the elements of \mathbb{V} would be a preferred interpretation of F' . Thus, $\text{prf}(F') = \mathbb{V}$ would be the case, which it is easy to see is actually false. Finally, we get $\Sigma_{\text{SFSADF}}^{\text{adm}} \not\subseteq \Sigma_{\text{AF}}^{\text{adm}}$ by observing that from *adm*(D) being realizable under *adm* in AFs it would follow that $\text{prf}(D)$ is realizable under *prf* in AFs. But we already showed that the latter is not the case. \square

Our final result on the admissibility-based semantics concerns the class of support-free ADFs (SFADFs). Recall that support-free symmetric ADFs (SFSADFs) have been defined in Section 3.3 in order to investigate subclasses of symmetric ADFs that satisfy certain properties.

Fig. 9. Expressiveness of subclasses of ADFs for $\sigma \in \{adm, prf, com\}$.

Theorem 24. For $\sigma \in \{prf, adm, com\}$, the following hold:

- $\Sigma_{SFSADF}^\sigma \subsetneq \Sigma_{SFADF}^\sigma$,
- $\Sigma_{SFADF}^\sigma \approx \Sigma_{ASSADF}^\sigma$,
- $\Sigma_{AF}^\sigma \subsetneq \Sigma_{SFADF}^\sigma$,
- $\Sigma_{SFADF}^\sigma \subsetneq \Sigma_{BADF}^\sigma$.

Proof. We separately prove the four relations provided in the theorem.

- Each SFSADF is a SFADF, that is, $\Sigma_{SFSADF}^\sigma \subseteq \Sigma_{SFADF}^\sigma$. To show that $\Sigma_{SFADF}^\sigma \not\subseteq \Sigma_{SFSADF}^\sigma$, let $\mathbb{V} = \{\{a \mapsto \mathbf{u}\}\}$. A witness of σ -realizability of \mathbb{V} in SFADFs is $D = (\{a\}, \{\varphi_a = \neg a\})$, for $\sigma \in \{prf, adm, com\}$. In contrast $\mathbb{V} \notin \Sigma_{SFSADF}^\sigma$.
- To show $\Sigma_{SFADF}^\sigma \not\subseteq \Sigma_{ASSADF}^\sigma$, let $\mathbb{V} = \{\{a \mapsto \mathbf{u}\}\}$, which is σ -realizable in SFADFs but not in ASSADFs. To show $\Sigma_{ASSADF}^\sigma \not\subseteq \Sigma_{SFADF}^\sigma$, let $\mathbb{V} = \{\{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}\}$. \mathbb{V} can be realized by $D = (\{a, b\}, \{\varphi_a = \neg b, \varphi_b = a\})$ in ASSADF under σ but not in SFADF.
- $\Sigma_{AF}^\sigma \subset \Sigma_{SFADF}^\sigma$ is clear. To show that $\Sigma_{SETAF}^\sigma \not\subseteq \Sigma_{AF}^\sigma$, let $\mathbb{V} = \{\{a \mapsto \mathbf{f}\}\}$ for $\sigma \in \{prf, com\}$, and respectively let $\mathbb{V}' = \{\{a \mapsto \mathbf{u}\}, \{a \mapsto \mathbf{f}\}\}$ for $\sigma = adm$. Both interpretation-sets are realizable by $D = (\{a\}, \{\varphi_a = \perp\})$ in SFADF under σ . In contrast $\mathbb{V} \notin \Sigma_{AF}^\sigma$.
- Each SFADF does not contain any dependent link; hence it is a BADF. The interpretations that are given in the proof of Theorem 21 work here to show that $\Sigma_{BADF}^\sigma \not\subseteq \Sigma_{SFADF}^\sigma$. \square

Our results comparing the expressivity of ASSADFs, SFSADFs, SFADFs, AFs, BADFs, and general ADFs w.r.t. the admissible, complete, and preferred semantics is summarised in Fig. 9. To complete the picture here we also incorporate results about the relative expressivity of AFs, BADFs, and general ADFs from [25].

The general picture for the stable semantics, which we proceed to investigate now, deviates somewhat from that of the admissibility based semantics. To start we remind the reader that stable models v, w of any ADF are always incomparable, i.e. $w^t \subseteq v^t$ implies $w^t = v^t$, see [31]. Now, in order to complete previous results from [32] comparing AFs, BADFs and general ADFs in terms of their expressivity w.r.t. the stable semantics, we make use of this fact and the forthcoming lemma, in order to prove that $\Sigma_{BADF}^{stb} \subseteq \Sigma_{SFSADF}^{stb}$.

Lemma 25. Any incomparable set of two-valued interpretations \mathbb{V} is stb-realizable in SFSADFs.

Proof. (sketch) To show that any incomparable set of two-valued interpretations is *stb*-realizable by some SFSADF, let \mathbb{V} be an incomparable set over arguments S (that is, each $v \in \mathbb{V}$ assigns \mathbf{t} or \mathbf{f} to every $s \in S$) and consider a SFSADF $D = (S, L, C)$ with the following acceptance conditions for $s \in S$:²

- If $v(s) = \mathbf{t}$ for every $v \in \mathbb{V}$ then $\varphi_s \equiv \top$.
- If $v(s) = \mathbf{f}$ for every $v \in \mathbb{V}$ then $\varphi_s \equiv \perp$.
- Otherwise, $\varphi_s = \bigvee_{v \in \mathbb{V}, v(s)=\mathbf{t}} \bigwedge_{v(t)=\mathbf{f} \wedge \exists w \in \mathbb{V}: (w(s)=\mathbf{f} \wedge w(t)=\mathbf{t})} \neg t$.

We show that D is a SFSADF and $stb(D) = \mathbb{V}$.

- By the definition of the acceptance condition of argument s in D , $s \notin par(s)$. Therefore, L is ir-reflexive. Further, $t \in par(s)$ iff $s \in par(t)$, that is, L is symmetric. In addition, all links in D are attacking. Thus, D is a SFSADF.
- To prove that $stb(D) = \mathbb{V}$, we show that $stb(D) \subseteq \mathbb{V}$ and also $\mathbb{V} \subseteq stb(D)$.
 - * To show $\mathbb{V} \subseteq stb(D)$, let $v \in \mathbb{V}$. We show that $v \in mod(D)$ and since SFSADFs are weakly coherent by Theorem 8, $v \in stb(D)$. Let $v(s) = \mathbf{t}$, we show that s is acceptable in v . (The proof for the case that $v(s) = \mathbf{f}$ is analogous). If s is assigned \mathbf{t} by each element of \mathbb{V} there is nothing to prove. Otherwise, there exists a $w \in \mathbb{V}$ s.t. $w(s) = \mathbf{f}$. Since \mathbb{V} is incomparable, there exists t such that $v(t) = \mathbf{f}$ and $w(t) = \mathbf{t}$. Therefore, $t \in par(s)$. The set of all arguments like t make a conjunctive clause of φ_s which guarantees that s is accepted in v .
 - * To show that $stb(D) \subseteq \mathbb{V}$, toward a contradiction, assume that there exists $v \in stb(D)$ such that $v \notin \mathbb{V}$. Since $stb(D)$ is incomparable, there exists $s \in S$ such that $v(s) = \mathbf{t}$ and $\varphi_s \not\equiv \top$. Further, the acceptance condition of s is not unsatisfiable, otherwise s has to be assigned to \mathbf{f} in v . Thus, there exists $v' \in \mathbb{V}$ in which s is assigned to \mathbf{t} . Let K be the set of all v' in which s is assigned to \mathbf{t} . Since $v \notin \mathbb{V}$ and $stb(D)$ is incomparable, in each $v' \in K$ there exists t such that $v(t) = \mathbf{t}$ and $v'(t) = \mathbf{f}$. Let T be a set of all arguments like t . It can be shown that either each conjunctive clause of φ_s contains a $t \in T$, or there exists a $t \in T$ such that each conjunctive clause of φ_t consists of an argument of T . The former means that s is deniable with respect to v , and the latter means that t is deniable in v . That is, v is not a two-valued model of D and since SFSADFs are weakly coherent, v is not a stable model of D . This is in contradiction with our assumption that there exists $v \in stb(D)$ such that $v \notin \mathbb{V}$. Therefore, $stb(D) \subseteq \mathbb{V}$. \square

Note that the incomparability condition for the set of two valued interpretations \mathbb{V} in the statement of Lemma 25 is necessary. For instance, $\mathbb{V} = \{\{a \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{f}\}\}$ is a set of two-valued interpretations, which are not incomparable, that is not *stb*-realizable in SFSADFs.

Theorem 26. $\Sigma_{AF}^{stb} \subsetneq \Sigma_{SFSADF}^{stb} \sim \Sigma_{SFADF}^{stb} \sim \Sigma_{ASSADF}^{stb} \sim \Sigma_{BADF}^{stb}$.

Proof. $\Sigma_{AF}^{stb} \subsetneq \Sigma_{BADF}^{stb}$ is shown in [32], and $\Sigma_{SFSADF}^{\sigma} \subseteq \Sigma_{ASSADF}^{\sigma} \subseteq \Sigma_{BADF}^{\sigma}$ and $\Sigma_{SFSADF}^{\sigma} \subseteq \Sigma_{SFADF}^{\sigma} \subseteq \Sigma_{BADF}^{\sigma}$, for each semantics σ , are clear. If we can show $\Sigma_{BADF}^{stb} \subseteq \Sigma_{SFSADF}^{stb}$ we are thus done. Let $\mathbb{V} \in \Sigma_{BADF}^{stb}$. Since \mathbb{V} is a set of incomparable two-valued interpretations, by Lemma 25, \mathbb{V} is *stb*-realizable in SFSADFs. \square

²This construction is a slight adaption of a result from [32].

In Example 6 we give an example of an interpretation-set that is *stb*-realizable in SFSADFs but not *stb*-realizable in AFs.

Example 6. Let $\mathbb{V} = \{\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}\}, \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}\}$ be an interpretation-set. A witness of *stb*-realizability of \mathbb{V} in SFSADFs is $D = (S, L, C)$ for which $S = \{a, b, c\}$ and the acceptance conditions are $\varphi_a = \neg b \vee \neg c$, $\varphi_b = \neg a \vee \neg c$ and $\varphi_c = \neg a \vee \neg b$. We show that \mathbb{V} is not *stb*-realizable in AFs. Towards a contradiction assume that there exists an AF $F = (A, R)$ such that $stb(F) = \mathbb{V}$. The set of arguments of F is S as these are the arguments occurring in $stb(F)$. The interpretations in \mathbb{V} imply that there is no link between a , b and c in F . Then, by the definition of the stable semantics for AFs, the interpretation $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{t}\}$ is also a stable interpretation of F . Therefore, $stb(F) \neq \mathbb{V}$. Thus, the interpretation-set \mathbb{V} is not *stb*-realizable in AFs and $\Sigma_{\text{SFSADF}}^{stb} \not\subseteq \Sigma_{\text{AF}}^{stb}$.

The final semantics to investigate is the model-semantics. We only need one technical lemma.

Lemma 27. For any SFADF $D = (S, L, C)$, $mod(D)$ is incomparable.

Proof. Toward a contradiction assume that there are $v, w \in mod(D)$ such that $v^{\mathbf{t}} \subset w^{\mathbf{t}}$. Let $B \subset S$ be the set of argument which are assigned to \mathbf{t} in w , but are assigned to \mathbf{f} by v . Moreover, let a be an argument which is denied in v and accepted in w . At least a parent of a has to be in B , otherwise, $\varphi_a^v = \varphi_a^w$. Assume $par(a) \cap B = \{b_1, \dots, b_n\}$. Since D is a SFADF, all links are attacking. Then, by the definition of attacking links $v|_{\mathbf{t}}^{b_1}(\varphi_a) = \mathbf{f}$. That is, a is denied with respect to $v_1 = v|_{\mathbf{t}}^{b_1}$. Following the same method construct $v_i = v|_{\mathbf{t}}^{b_i}$, for $1 \leq i \leq n$. It is obvious that a is denied in each v_i , in particular, a is denied in v_n , which is equal to w . This is a contradiction to the assumption that a is accepted in w . Therefore, $mod(D)$ is a set of incomparable interpretations. \square

Theorem 28. $\Sigma_{\text{AF}}^{mod} \subsetneq \Sigma_{\text{SFSADF}}^{mod} \sim \Sigma_{\text{SFADF}}^{mod} \sim \Sigma_{\text{ASSADF}}^{mod} \subsetneq \Sigma_{\text{BADF}}^{mod}$.

Proof. We first argue that $\Sigma_{\mathcal{F}}^{mod} = \Sigma_{\mathcal{F}}^{stb}$ for $\mathcal{F} \in \{\text{SFSADF}, \text{SFADF}, \text{ASSADF}\}$. This follows immediately for $\mathcal{F} \in \{\text{SFSADF}, \text{ASSADF}\}$ since both ASSADFs and SFSADFs are weakly coherent (cf. Theorem 5 and Theorem 8). To show the relation for $\mathcal{F} = \text{SFADF}$ we need to show $\Sigma_{\text{SFADF}}^{mod} \subseteq \Sigma_{\text{SFADF}}^{stb}$ (the other inclusion follows by standard properties of semantics). Let $\mathbb{V} \in \Sigma_{\text{SFADF}}^{mod}$. By Lemma 27, \mathbb{V} is incomparable and by Lemma 25, any incomparable set of two-valued interpretations is *stb*-realizable in SFSADF. Therefore, there is a SFSADF D' such that $stb(D') = \mathbb{V}$. Since each SFSADF is also an SFADF, $\mathbb{V} \in \Sigma_{\text{SFADF}}^{stb}$.

Using these relations and observing that *stb* and *mod* are equivalent for AFs, it follows from Theorem 26 that $\Sigma_{\text{AF}}^{mod} \subsetneq \Sigma_{\text{SFSADF}}^{mod} \sim \Sigma_{\text{SFADF}}^{mod} \sim \Sigma_{\text{ASSADF}}^{mod}$ holds. Finally, $\Sigma_{\text{ASSADF}}^{mod} \subsetneq \Sigma_{\text{BADF}}^{mod}$ is by Theorem 21. \square

Figure 10 summarises the results regarding expressivity w.r.t. the model and stable semantics expressed in Theorem 26 and Theorem 28. Again, to complete the picture we make use of results from [32] ($\Sigma_{\text{BADF}}^{stb} \sim \Sigma_{\text{ADF}}^{stb}$ and $\Sigma_{\text{BADF}}^{mod} \subset \Sigma_{\text{ADF}}^{mod}$).

We conclude this section by pointing out that the realisability relationships depicted in Fig. 10 change when we restrict the cardinality of the interpretation sets. As it turns out, any set of interpretations of size 2 obtained from an ADF when evaluated using the stable semantics is also realizable in AFs.




Fig. 10. Expressiveness AFs, SFSADFs, SFADFs, ASSADFs, BADFs, ADFs for $\sigma \in \{\text{stb}, \text{mod}\}$.

Proposition 29. *Suppose that $|\mathbb{V}| = 2$ and \mathbb{V} is *stb-realizable* in ADFs. Then \mathbb{V} is *stb-realizable* in AFs.*

Proof. Let $\mathbb{V} = \{v_1, v_2\}$ be a set of interpretations that is *stb-realizable* in ADFs, i.e. there exists an ADF $D = (S, L, C)$ such that $\text{stb}(D) = \mathbb{V}$. Construct an AF $F = (A, R)$ by setting $A = S$ and $R = \{(a, b) \mid v_i(a) = \mathbf{t}, v_i(b) = \mathbf{f}, v_j(a) = \mathbf{f}, 1 \leq i \neq j \leq 2\}$. To prove that $\text{stb}(F) = \mathbb{V}$, take $v_i \in \mathbb{V}$. First, there is no attack between arguments with $v_i(a) = \mathbf{t}$. Moreover, if $v_i(b) = \mathbf{f}$ then, since neither $v_1 \leq_i v_2$ nor $v_2 \leq_i v_1$, there must be some $a \in A$ with $v_i(a) = \mathbf{t}$ and $v_j(a) = \mathbf{f}$, hence $(a, b) \in R$. Therefore v_i is a stable interpretation of F . That is, \mathbb{V} is *stb-realizable* in AFs. \square

5. Discussion

Motivated by crucial results on the semantic properties of acyclic (or well founded) [13], symmetric [10], and odd-length-cycle-free [15] AFs, in this work we investigated analogous classes for ADFs and their properties. We showed that for acyclic ADFs, just as is the case for acyclic AFs, the different semantics coincide. On the other hand, we demonstrated that the properties of coherence and relatively-groundedness that hold for symmetric AFs do not carry over to symmetric ADFs. The latter impelled us to go on a quest for an appropriate subclass of symmetric ADFs for which some form of coherence holds. In the process we defined several subclasses, in particular acyclic support symmetric ADFs (ASSADFs) and support-free symmetric ADFs (SFSADFs) which we show satisfy a weaker form of coherence.

Also odd-length-cycle-free ADFs do not satisfy coherence (which is the case for AFs), but here we were able to show that this property does hold for SFSADFs. In fact this is the case for a superclass of SFSADFs (which also contains AFs with collective attacks), which we dubbed SFADFs. This property also allowed us to distinguish ASSADFs from SFSADFs in that odd-length-cycle-free ASSADFs are not coherent in general.

The motivation behind this line of investigation lies in the fact that different semantics show different complexities [35]. It is thus valuable to know under which circumstances higher complexities can be avoided. Acyclicity is a positive example since the coincidence with grounded semantics shows that, for instance, the more complex preferred semantics becomes easier for this class of frameworks. The practical implication is as follows: an ADF solver should check for acyclicity before computing preferred interpretations, since in case the ADF to be treated is acyclic, the easier procedure for grounded semantics suffices to do the job. As we have shown, such a strategy does not carry over to symmetric ADFs. This is in contrast to symmetric AFs where coherence holds, i.e. the (more complex) preferred semantics coincides with the (easier) stable semantics. Nonetheless, there is still a chance that symmetric ADFs are of practical help. In contrast to acyclic ADFs where the complexity drop is immediate, our

results underline that dedicated complexity analyses for symmetric ADFs should be considered as future work.

As a further contribution and also following in the footsteps of work for AFs [16], we considered the subclasses of ADFs we introduced (ASSADFs, SFADFs, and SFSADFs) in terms of their expressivity as can be gleaned from their signatures. Here SFSADFs are a strict subset of ASSADFs for the admissibility-based semantics we considered, while SFADFs are incomparable w.r.t. ASSADFs (and a strict superset of SFSADFs). On the other hand the signatures of ASSADFs, SFADFs, and SFSADFs coincide for the model and stable semantics.

We also complemented previous work on expressivity of AFs and ADFs [25,29,33] by comparing the expressivity of ASSADFs, SFADFs, and SFSADFs with that of AFs, bipolar ADFs (BADFs), and ADFs. Here ASSADFs and SFSADFs are incomparable with AFs for the admissibility semantics, while SFADFs are strictly more expressive. ASSADFs, SFADFs, and SFSADFs are strictly more expressive for the model and stable semantics. On the other hand, they are strictly less expressive than BADFs for the model and admissibility based semantics, while they coincide in expressivity with BADFs and general ADFs for the stable semantics.

This work is an elaboration on the more theoretical aspects of our work presented in [11]. There we had also included results on an empirical evaluation of some of the main systems for ADFs (QADF [12], YADF [3], and `gODIAMOND` [34]) on acyclic vs. non-acyclic ADFs. These show usually only a slight improvement of these systems (which do not detect subclasses of ADFs) on the acyclic instances despite the fact that the results we present in this work indicate that even the most difficult of reasoning problems become tractable for acyclic ADFs. Thus, our work offers further guidelines for designing more efficient systems for ADFs. As a notable example and as we suggested in the introduction, backdoor approaches that utilize the distance to subclasses with easier complexity, such as the systems `cegartix` [19] for AFs as well as the very recent `k++ADF` [24] for ADFs, can be devised on the basis of our results.

Finally, also from a theoretical perspective our work leaves open several further avenues to explore. In particular, we consider to extend our studies to other ADF semantics [22,28] as well as generalizations of ADFs [7].

Acknowledgements

This research has been supported by FWF projects I2854, P30168, and W1255, and by DFG project 389792660 – TRR 248. The second researcher is currently embedded in the Center of Data Science & Systems Complexity (DSSC) Doctoral Programme at the University of Groningen.

References

- [1] L. Al-Abdulkarim, K. Atkinson and T.J.M. Bench-Capon, A methodology for designing systems to reason with legal cases using abstract dialectical frameworks, *Artificial Intelligence Law* **24**(1) (2016), 1–49. doi:[10.1007/s10506-016-9178-1](https://doi.org/10.1007/s10506-016-9178-1).
- [2] P. Baroni, M. Caminada and M. Giacomin, An introduction to argumentation semantics, *Knowledge Eng. Review* **26**(4) (2011), 365–410. doi:[10.1017/S0269888911000166](https://doi.org/10.1017/S0269888911000166).
- [3] G. Brewka, M. Diller, G. Heissenberger, T. Linsbichler and S. Woltran, Solving advanced argumentation problems with answer-set programming, in: *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*, AAAI Press, 2017, pp. 1077–1083.
- [4] G. Brewka, S. Ellmauthaler, H. Strass, J.P. Wallner and S. Woltran, Abstract dialectical frameworks revisited, in: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, IJCAI/AAAI, 2013, pp. 803–809.

- [5] G. Brewka, S. Ellmauthaler, H. Strass, J.P. Wallner and S. Woltran, Abstract dialectical frameworks. An overview, *IfCoLog Journal of Logics and their Applications* **4**(8) (2017), 2263–2318.
- [6] G. Brewka, S. Ellmauthaler, H. Strass, J.P. Wallner and S. Woltran, Abstract dialectical frameworks: An overview, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018.
- [7] G. Brewka, H. Strass, J.P. Wallner and S. Woltran, Weighted abstract dialectical frameworks, in: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI, 2018)*, AAAI Press, 2018, pp. 1779–1786.
- [8] G. Brewka and S. Woltran, Abstract dialectical frameworks, in: *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR 2010)*, AAAI Press, 2010, pp. 102–111.
- [9] E. Cabrio and S. Villata, Abstract dialectical frameworks for text exploration, in: *Proceedings of International Conference on Agents and Artificial Intelligence (ICAART 2016)*, SciTePress, 2016, pp. 85–95.
- [10] S. Coste-Marquis, C. Devred and P. Marquis, Symmetric argumentation frameworks, in: *Proceedings of the 8th European Conference On Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2005)*, Springer, 2005, pp. 317–328. doi:[10.1007/11518655_28](https://doi.org/10.1007/11518655_28).
- [11] M. Diller, A. Keshavarzi Zafarghandi, T. Linsbichler and S. Woltran, Investigating subclasses of abstract dialectical frameworks, in: *Proceedings of Computational Models of Argument (COMMA 2018)*, IOS Press, 2018, pp. 61–72.
- [12] M. Diller, J.P. Wallner and S. Woltran, Reasoning in abstract dialectical frameworks using quantified Boolean formulas, in: *Proceedings of Computational Models of Argument (COMMA 2014)*, IOS Press, 2014, pp. 241–252.
- [13] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artificial Intelligence* **77**(2) (1995), 321–357. doi:[10.1016/0004-3702\(94\)00041-X](https://doi.org/10.1016/0004-3702(94)00041-X).
- [14] P.E. Dunne, Computational properties of argument systems satisfying graph-theoretic constraints, *Artificial Intelligence* **171**(10–15) (2007), 701–729. doi:[10.1016/j.artint.2007.03.006](https://doi.org/10.1016/j.artint.2007.03.006).
- [15] P.E. Dunne and T.J.M. Bench-Capon, Coherence in finite argument systems, *Artificial Intelligence* **141**(1/2) (2002), 187–203. doi:[10.1016/S0004-3702\(02\)00261-8](https://doi.org/10.1016/S0004-3702(02)00261-8).
- [16] P.E. Dunne, W. Dvořák, T. Linsbichler and S. Woltran, Characteristics of multiple viewpoints in abstract argumentation, *Artificial Intelligence* **228** (2015), 153–178. doi:[10.1016/j.artint.2015.07.006](https://doi.org/10.1016/j.artint.2015.07.006).
- [17] W. Dvořák and P.E. Dunne, Computational problems in formal argumentation and their complexity, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, pp. 631–687.
- [18] W. Dvořák, J. Fandinno and S. Woltran, On the expressive power of collective attacks, *Argument & Computation* **10**(2) (2019), 191–230. doi:[10.3233/AAC-190457](https://doi.org/10.3233/AAC-190457).
- [19] W. Dvořák, M. Järvisalo, J.P. Wallner and S. Woltran, Complexity-sensitive decision procedures for abstract argumentation, *Artificial Intelligence* **206** (2014), 53–78. doi:[10.1016/j.artint.2013.10.001](https://doi.org/10.1016/j.artint.2013.10.001).
- [20] W. Dvořák, S. Ordyniak and S. Szeider, Augmenting tractable fragments of abstract argumentation, *Artificial Intelligence* **186** (2012), 157–173. doi:[10.1016/j.artint.2012.03.002](https://doi.org/10.1016/j.artint.2012.03.002).
- [21] G. Flouris and A. Bikakis, A comprehensive study of argumentation frameworks with sets of attacking arguments, *International Journal of Approximate Reasoning* **109** (2019), 55–86. doi:[10.1016/j.ijar.2019.03.006](https://doi.org/10.1016/j.ijar.2019.03.006).
- [22] S.A. Gaggl, S. Rudolph and H. Strass, On the computational complexity of naive-based semantics for abstract dialectical frameworks, in: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, IJCAI/AAAI, 2015, pp. 2985–2991.
- [23] A. Keshavarzi Zafarghandi, Investigating Subclasses of Abstract Dialectical Frameworks, Master’s thesis, TU Wien, 2017.
- [24] T. Linsbichler, M. Maratea, A. Niskanen, J.P. Wallner and S. Woltran, Novel algorithms for abstract dialectical frameworks based on complexity analysis of subclasses and SAT solving, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018)*, Ijcai.org, 2018, pp. 1905–1911.
- [25] T. Linsbichler, J. Pührer and H. Strass, A uniform account of realizability in abstract argumentation, in: *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*, IOS Press, 2016, pp. 252–260.
- [26] D. Neugebauer, Generating defeasible knowledge bases from real-world argumentations using D-BAS, in: *Proceedings of the 1st Workshop on Advances in Argumentation in Artificial Intelligence Co-Located with XVI International Conference of the Italian Association for Artificial Intelligence*, CEUR-WS.org, 2017, pp. 105–110.
- [27] S.H. Nielsen and S. Parsons, A generalization of Dung’s abstract framework for argumentation: Arguing with sets of attacking arguments, in: *Argumentation in Multi-Agent Systems (ArgMAS 2006)*, Springer, 2006, pp. 54–73.
- [28] S. Polberg, Understanding the abstract dialectical framework, in: *Logics in Artificial Intelligence – 15th European Conference (JELIA 2016)*, Springer, 2016, pp. 430–446. doi:[10.1007/978-3-319-48758-8_28](https://doi.org/10.1007/978-3-319-48758-8_28).
- [29] J. Pührer, Realizability of three-valued semantics for abstract dialectical frameworks, in: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, AAAI Press, 2015, pp. 3171–3177.
- [30] J. Pührer, ArgueApply: A mobile app for argumentation, in: *Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2017)*, Springer, 2017, pp. 250–262. doi:[10.1007/978-3-319-61660-5_23](https://doi.org/10.1007/978-3-319-61660-5_23).

- [31] H. Strass, Approximating operators and semantics for abstract dialectical frameworks, *Artificial Intelligence* **205** (2013), 39–70. doi:[10.1016/j.artint.2013.09.004](https://doi.org/10.1016/j.artint.2013.09.004).
- [32] H. Strass, Expressiveness of two-valued semantics for abstract dialectical frameworks, *Journal of Artificial Intelligence Research* **54** (2015), 193–231. doi:[10.1613/jair.4879](https://doi.org/10.1613/jair.4879).
- [33] H. Strass, Expressiveness of two-valued semantics for abstract dialectical frameworks, *Journal of Artificial Intelligence Research* **54** (2015), 193–231. doi:[10.1613/jair.4879](https://doi.org/10.1613/jair.4879).
- [34] H. Strass and S. Ellmauthaler, *GoDIAMOND 0.6.6 – ICCMA 2017 System Description*, 2017, available at: <http://argumentationcompetition.org/2017/goDIAMOND.pdf>.
- [35] H. Strass and J.P. Wallner, Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory, *Artificial Intelligence* **226** (2015), 34–74. doi:[10.1016/j.artint.2015.05.003](https://doi.org/10.1016/j.artint.2015.05.003).
- [36] J.P. Wallner, Structural constraints for dynamic operators in abstract argumentation, in: *Proceedings of Computational Models of Argument (COMMA 2018)*, IOS Press, 2018, pp. 73–84.