Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Realizability of Three-Valued Semantics for Abstract Dialectical Frameworks*

Jörg Pührer
Leipzig University
Leipzig, Germany
puehrer@informatik.uni-leipzig.de

Abstract

We investigate fundamental properties of three-valued semantics for abstract dialectical frameworks (ADFs). In particular, we deal with *realizability*, i.e., the question whether there exists an ADF that has a given set of interpretations as its semantics. We provide necessary and sufficient conditions that hold for a set of three-valued interpretations whenever there is an ADF realizing it under admissible, complete, grounded, or preferred semantics. Moreover, we discuss how to construct such an ADF in case of realizability. Our results lay the ground for studying the expressiveness of ADFs under three-valued semantics. As a first application we study implications of our results on the existence of certain join operators on ADFs.

1 Introduction

Abstract dialectical frameworks (ADFs) [Brewka and Woltran, 2010] have been introduced as a generalization of Dung's argumentation frameworks (AFs) [Dung, 1995]. In contrast to AFs, where relations between arguments are restricted to bilateral attacks, the acceptance of a statement in an ADF may depend on an arbitrary function of the acceptance of other statements. This way, it allows for expressing more complex dependencies between arguments. Traditionally, the semantics of AFs is based on so-called *extensions*, i.e., sets of arguments that are in some sense compatible with each other. Extensions can also be seen as dedicated (two-valued) interpretations of propositional logic where each argument stands for a propositional variable. Caminada and Gabbay [2009] introduced a labeling based semantics for AFs, where every argument is assigned one of the labels `in`, `out`, or `undec` with respect to a labeling. Thus, a labeling can be seen as a three-valued interpretation. When an argument is `in` or `out` it means that it is accepted or rejected, respectively. The third label `undec` allows for expressing a point of view without an explicit judgment on the acceptance or rejection of the argument. Brewka *et al.* [2013] introduced various three-valued semantics also for ADFs that generalize corresponding labeling semantics for AFs. In this work, our aim is to characterize these semantics. In particular, we deal with the question whether a given set V of three-valued interpretations is *realizable* by an ADF, i.e., whether there is an ADF that has V as its set of admissible, complete, grounded, respectively, preferred interpretations. We address the problem by providing necessary and sufficient conditions that hold for V whenever there is an ADF realizing V under these semantics. In doing so, we continue a series of works on signatures of argumentation semantics [Dunne *et al.*, 2014; Baumann *et al.*, 2014; Dyrkolbotn, 2014; Strass, 2015]. This important line of research fosters the understanding of argumentation formalisms and is a prerequisite for studying diverse aspects of their expressiveness (cf. [Gogic *et al.*, 1995; Dunne *et al.*, 2014; Baumann *et al.*, 2014; Strass, 2015; Eiter *et al.*, 2013]). With suitable realizability characterizations at hand we can, e.g., compare what different formalisms (or one formalism under different semantics) can express, show whether a language is suitable to capture a given problem, or identify expressive subclasses of a formalism.

While other previous works deal with AFs, Strass [2015] also addressed ADFs, investigating realizability under two-valued semantics, in particular stable and model semantics [Brewka *et al.*, 2013]. We close a gap by characterizing the three-valued semantics introduced by Brewka *et al.* For admissible semantics we identify interpretations that must be admissible if all interpretations in a given set V are considered admissible. It turns out that V is realizable exactly when it already contains these induced interpretations. In the case of grounded and preferred semantics the characterization depends on simple syntactic conditions on V. For complete semantics, we reduce the existence of a realizing ADF to the existence of a characterization function and provide a complexity argument in favor of this solution. In addition, we provide a couple of simpler conditions that are necessary but not sufficient for complete realizability. Our results are based on constructive proofs and we discuss how to obtain an ADF realizing a set V whenever V is realizable.

As a first application of our characterizations we analyze whether there can be a join operator for ADFs that adheres to a certain desired requirement. Here, we get different results for different semantics, e.g., while the answer is positive for admissible semantics, we show that no such operator can exist for complete semantics.

*The German Research Foundation (DFG) supported this work under grant BR-1817/7-1.

2 Preliminaries

Unless stated otherwise, we assume that S is a fixed non-empty and finite set of statements. Intuitively, a statement can be seen as a position in a conversation and corresponds to an argument in the context of AFs. Following Brewka and Woltran [2010], an ADF is a tuple $\langle S, L, C\rangle$, where S is the set of statements and $L \subseteq S^2$ is a set of links between statements. The idea is that the relation L determines for each statement s on which other statements its acceptance might depend. We call such a statement a *parent* of s and denote the set of parents of s by $par(s) = \{s' \mid \langle s', s\rangle \in L\}$. The final component $C = \{C_s \mid s \in S\}$ of an ADF is a set of total functions $C_s : 2^{par(s)} \to \{\mathbf{t}, \mathbf{f}\}$, one for each statement s. C_s is called acceptance condition of s and if it maps some $R \subseteq par(s)$ to $\mathbf{t}$ that means that s is accepted if all statements in R are accepted and all statements in $par(s) \setminus R$ are rejected. We denote the set of all ADFs by $\mathcal{D}$.

For convenience, we will use a more compact notation for ADFs, in which the parent relation L is implicit and acceptance conditions are represented by propositional formulas (acceptance formulas): We represent an ADF as above as a set D that contains one pair $\langle s, \varphi\rangle$ for each $s \in S$ where φ is a propositional formula over S such that for each $R \subseteq S$, $C_s(R) = \mathbf{t}$ iff $R \models \varphi$ (where R is viewed as a two-valued propositional interpretation and $\models$ is defined as usual).

We define semantics for ADFs as in Brewka *et al.* [2013] using a slightly different notation. A (three-valued) interpretation v is a function $v : S \to \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$, assigning every statement one of the truth values true ($\mathbf{t}$), false ($\mathbf{f}$), or undefined ($\mathbf{u}$). Next, we introduce some notation that helps us to work with interpretations. For an interpretation v we use the sets $\mathrm{T}_v = \{s \in S \mid v(s) = \mathbf{t}\}$, $\mathrm{F}_v = \{s \in S \mid v(s) = \mathbf{f}\}$, and $\mathrm{U}_v = \{s \in S \mid v(s) = \mathbf{u}\}$ to collect the statements with the same truth values under v. $\mathrm{TF}_v = \mathrm{T}_v \cup \mathrm{F}_v$ is the set of *defined* statements under v and $\mathcal{V}$ denotes the set of all interpretations. We sometimes represent interpretations by listing their defined statements, where statements mapped to $\mathbf{f}$ are barred, e.g., the interpretation v with $\mathrm{T}_v = \{a\}$, $\mathrm{F}_v = \{b\}$, $\mathrm{U}_v = \{c\}$ is denoted by $a\bar{b}$. Furthermore, the interpretation that maps all statements to $\mathbf{u}$ is denoted by $v_{\mathbf{u}}$. The *information ordering* $\leq_i$ between interpretations is defined as $\leq_i = \{\langle v, v'\rangle \in \mathcal{V}^2 \mid \mathrm{T}_v \subseteq \mathrm{T}_{v'}, \mathrm{F}_v \subseteq \mathrm{F}_{v'}\}$. Thus, for two different interpretations v and v', $v \leq_i v'$ means that v' extends v by assigning more statements a classical truth value. We write $v <_i v'$ if $v \leq_i v'$ and $v \neq v'$. The semantics for ADFs we are interested in are defined via two-valued extensions of three-valued interpretations with respect to the information ordering. Hence, we identify three-valued interpretations that do not assign $\mathbf{u}$ to any statement with its corresponding two-valued interpretation. Such an interpretation assigns a truth value to a propositional formula in the standard way. By $\mathcal{V}_2 = \{v \in \mathcal{V} \mid \mathrm{U}_v = \emptyset\}$ we denote the set of two-valued interpretations. Given an interpretation v, the set $[v]_2 = \{v_2 \in \mathcal{V}_2 \mid v \leq_i v_2\}$ collects the two-valued extensions of v.

The central device for defining three-valued semantics for ADFs is the *consensus operator* that maps an interpretation v to another interpretation v', where, intuitively, v' maps statements to truth values of their acceptance formulas on which the two-valued extensions of v agree: Given an ADF D, the consensus for an interpretation v is the interpretation $\Gamma_D(v)$, where $\mathrm{T}_{\Gamma_D(v)} = \{s \mid \langle s, \varphi\rangle \in D, v'(\varphi) = \mathbf{t}$ for every $v' \in [v]_2\}$ and $\mathrm{F}_{\Gamma_D(v)} = \{s \mid \langle s, \varphi\rangle \in D, v'(\varphi) = \mathbf{f}$ for every $v' \in [v]_2\}$.

Definition 1 ([Brewka *et al.*, 2013]) *Let D be an ADF. An interpretation v is*

- admissible *for D if $v \leq_i \Gamma_D(v)$;*
- complete *for D if $v = \Gamma_D(v)$;*
- grounded *for D if v is complete for D and every $v' \in \mathcal{V}$ with $v' <_i v$ is not complete for D;*
- preferred *for D if v is admissible for D and every $v' \in \mathcal{V}$ with $v <_i v'$ is not admissible for D.*

By $\mathrm{adm}(D)$, $\mathrm{com}(D)$, $\mathrm{grd}(D)$, *respectively,* $\mathrm{pre}(D)$, *we denote the set of admissible, complete, grounded, respectively, preferred interpretations for D.*

It is known that every grounded and every preferred interpretation for a given ADF D is complete for D and that every complete interpretation is admissible for D [Brewka *et al.*, 2013]. For intuition on the different types of semantics we refer to Baroni *et al.* [2011] who discuss them for AFs. Due to space limitations we present only selected proofs.

3 Characterizing Three-Valued Semantics

Before analyzing concrete semantics, we want to start with a general observation about the nature of three-valued semantics for ADFs. That is, there is no three-valued semantics under which every set of interpretations can be realized. More precisely, there is no function $\sigma : \mathcal{D} \to 2^{\mathcal{V}}$ assigning a set of interpretations to an ADF such that for every $V \subseteq \mathcal{V}$ there is some ADF $D \in \mathcal{D}$ with $\sigma(F) = V$. This is due to a simple cardinality argument: there are more sets of three-valued interpretations than ADFs. There are 2^{2^n} Boolean functions in n variables and, in order to build an ADF, one of these functions is required for each statement. As a consequence, the number of ADFs with n statements is 2^{n2^n}. On the other hand, for n statements, there are 3^n three-valued interpretations. Therefore, 2^{3^n} sets of interpretations exist. The claim holds as for every natural number n we have $2^{n2^n} < 2^{3^n}$.

Approaching from this negative result we now turn our attention to the question which sets *can* be realized under different semantics and how. As a teaser, we want to provide some numbers for ADFs with three statements on a side note: There are 16,777,216 ADFs with three statements and 134,217,728 corresponding sets of interpretations. From the latter, only 77,712 can be realized under admissible semantics, 16,618 under complete semantics, 2088 under preferred semantics, and, finally, only 27 under grounded semantics.

3.1 Admissible Semantics

In this section, we want to characterize sets V of interpretations that constitute the admissible interpretations of an ADF. We start with the observation that the admissibility of interpretations of an ADF require other interpretations to be admissible as well. For identifying them, we introduce the following notion that depends on a set of interpretations.

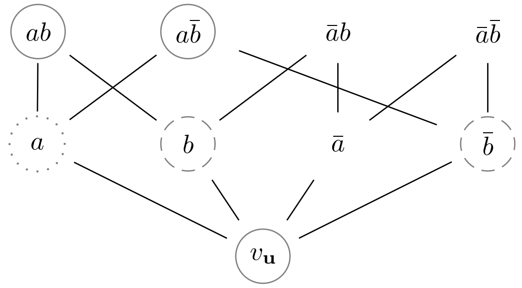

Figure 1: All interpretations for $S = \{a, b\}$ arranged in the lattice $\langle \mathcal{V}, \leq_i \rangle$. Circled nodes represent interpretations from a given set V, where dotted circles express membership in $V^{\text{def}}_{\text{adm}}(a)$ and dashed circles in $V^{\text{def}}_{\text{adm}}(b)$. V is adm-closed.

Definition 2 *Let $V \subseteq \mathcal{V}$ be a set of interpretations and $v \in \mathcal{V}$ an interpretation. We call v* adm-induced *by V if for every $s \in \mathrm{TF}_v$ and every $v_2 \in [v]_2$ there is some $v' \in V$ such that $v' \leq_i v_2$ and $v'(s) = v(s)$.*

$V^{\text{ind}}_{\text{adm}}$ denotes the set of interpretations adm-induced by V.

Note that the definition implies that $V \subseteq V^{\text{ind}}_{\text{adm}}$; hence every element of V is adm-induced by V, but the converse does not hold for arbitrary sets of interpretations. If it does we call V *adm-closed*:

Definition 3 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. V is* adm-closed *if $V^{\text{ind}}_{\text{adm}} \subseteq V$.*

It turns out that adm-closedness is a necessary condition for realizability under admissible semantics.

Proposition 1 *For every ADF D, $\mathrm{adm}(D)$ is adm-closed.*

Proof. Consider some interpretation $v \in \mathrm{adm}(D)^{\text{ind}}_{\text{adm}}$. It suffices to show that $v \in \mathrm{adm}(D)$, i.e. that for all $\langle s, \varphi \rangle \in D$ where $s \in \mathrm{TF}_v$ and all $v_2 \in [v]_2$ it holds that $v_2(\varphi) = v(s)$. Consider some arbitrary $\langle s, \varphi \rangle \in D$ where $s \in \mathrm{TF}_v$ and some $v_2 \in [v]_2$. From $v \in \mathrm{adm}(D)^{\text{ind}}_{\text{adm}}$, we get that there is some $v' \in \mathrm{adm}(D)$ with $v'(s) = v(s)$ and $v' \leq_i v_2$. But then, from $v' \in \mathrm{adm}(D)$ and $v_2 \in [v']_2$, we get that $v_2(\varphi) = v'(s)$. Thus, also $v_2(\varphi) = v(s)$. □

Next, we show constructively that adm-closedness is also a sufficient condition for realizability. As we will see, if V is realizable under admissible semantics, only a particular subset of V is sufficient to determine all admissible interpretations. This subset is formed by the adm-defining interpretations of each statement as defined next.

Definition 4 *Given $V \subseteq \mathcal{V}$ and $s \in S$, the* set of adm-defining interpretations *of s in V is*

$$V^{\text{def}}_{\text{adm}}(s) = \{v \in V \mid s \in \mathrm{TF}_v \cap \bigcap_{v' \in V, v' <_i v} \mathrm{U}_{v'}\}.$$

As depicted in Figure 1, this set consists of the $\leq_i$-minimal interpretations in which a Boolean value is assigned to s. Based on that notion, we can define acceptance formulas for statements that allow us to construct a canonical ADF for a given set of interpretations.

Definition 5 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. The* canonical acceptance formula *of a statement $s \in S$ in V is given by $\psi_V(s) = \psi^+_V(s) \vee \psi^-_V(s)$ where*

$$\psi^+_V(s) = \bigvee_{v \in V^{\text{def}}_{\text{adm}}(s), s \in \mathrm{T}_v} \Big(\bigwedge_{s' \in \mathrm{T}_v} s' \wedge \bigwedge_{s' \in \mathrm{F}_v} \neg s' \Big) \quad \textit{and}$$

$$\psi^-_V(s) = \neg s \wedge \bigwedge_{v \in V^{\text{def}}_{\text{adm}}(s), s \in \mathrm{F}_v} \Big(\bigvee_{s' \in \mathrm{T}_v} \neg s' \vee \bigvee_{s' \in \mathrm{F}_v} s' \Big)$$

The canonical ADF *of V is given by*

$$\mathcal{CAN}(V) = \{\langle s, \psi_V(s) \rangle \mid s \in S\}.$$

Note that $\mathcal{CAN}(V)$ is a well-defined ADF.

Proposition 2 *Let V be an adm-closed set of interpretations. Then, $V = \mathrm{adm}(\mathcal{CAN}(V))$.*

Proof. ($\subseteq$) Towards a contradiction, consider some interpretation $v \in V$ such that $v \notin \mathrm{adm}(\mathcal{CAN}(V))$. From the latter we conclude that there is some $s \in \mathrm{TF}_v$ and some $v_2 \in [v]_2$ such that $v_2(\psi_V(s)) \neq v(s)$. Moreover, as $v \in V$ and $s \in \mathrm{TF}_v$ there must be some $v' \in V^{\text{def}}_{\text{adm}}(s)$ with $v' \leq_i v \leq_i v_2$ with $v'(s) = v(s)$. Consider the case that $s \in \mathrm{T}_v$. Then, from $v_2(\psi_V(s)) \neq v(s)$ follows that $v_2(\psi^+_V(s)) = \mathbf{f}$. Consequently, also

$$v_2\Big(\bigwedge_{s' \in \mathrm{T}_{v'}} s' \wedge \bigwedge_{s' \in \mathrm{F}_{v'}} \neg s' \Big) = \mathbf{f}.$$

Hence, there is some $s' \in \mathrm{TF}_{v'}$ such that $v'(s') \neq v_2(s')$, being a contradiction to $v' \leq_i v_2$. For the case that $s \in \mathrm{F}_v$, we have $v_2(\psi^-_V(s)) = \mathbf{t}$ and, in particular,

$$v_2\Big(\bigvee_{s' \in \mathrm{T}_{v'}} \neg s' \vee \bigvee_{s' \in \mathrm{F}_{v'}} s' \Big) = \mathbf{t}.$$

Again, there must be some $s' \in \mathrm{TF}_{v'}$ such that $v'(s') \neq v_2(s')$ which contradicts $v' \leq_i v_2$.

($\supseteq$) Now assume $v \in \mathrm{adm}(\mathcal{CAN}(V))$ but $v \notin V$. As V is adm-closed, v is not adm-induced by V. As a consequence, there is some $s \in \mathrm{TF}_V$ and some $v_2 \in [v]_2$ such that ($\star$) for each $v' \in V^{\text{def}}_{\text{adm}}(s)$ it cannot hold that $v' \leq_i v_2$. From $v \in \mathrm{adm}(\mathcal{CAN}(V))$ we get that $v_2(\psi_V(s)) = v(s)$. Consider the case that $s \in \mathrm{T}_v$. Then, from $v_2(\psi_V(s)) = \mathbf{t}$ follows that $v_2(\psi^+_V(s)) = \mathbf{t}$ as we know that $s \in \mathrm{T}_v$ implies $v_2(\psi^-_V(s)) = \mathbf{f}$. Hence, by definition of $\psi^+_V(s)$, there is some $v' \in V^{\text{def}}_{\text{adm}}(s)$ such that $v' \leq_i v_2$. This is a contradiction to ($\star$). In case that $s \in \mathrm{F}_v$ it must hold that $v_2(\psi^-_V(s)) = \mathbf{f}$. Again, this means that there is some $v' \in V^{\text{def}}_{\text{adm}}(s)$ such that $v' \leq_i v_2$, leading to a contradiction to ($\star$). □

We are ready to characterize realizability for admissible semantics.

Theorem 1 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. There is an ADF D such that $\mathrm{adm}(D) = V$ iff V is adm-closed.*

Proof. The 'if'-direction follows from Proposition 2 and the 'only if'-direction from Proposition 1. □

While adm-closedness characterizes sets of admissible interpretations, checking the property for a given set V is expensive: we need to check for every interpretation v that is not contained in V whether it is adm-induced by V. Following the definition of being adm-induced, this would require to find a $v' \in V$ with $v' \leq_i v_2$ and $v'(s) = v(s)$ for every two-valued extension v_2 of v and every $s \in \mathrm{TF}_v$.

In the remainder of the section, we show how we can narrow the search for such a v'. To this end, we need the following notions.

Definition 6 *Let $v, v' \in \mathcal{V}$ be interpretations. Then,*

- *v and v' are* compatible *if $(\mathrm{T}_v \cap \mathrm{F}_{v'}) \cup (\mathrm{T}_{v'} \cap \mathrm{F}_v) = \emptyset$;*
- *if v and v' are compatible we denote the interpretation v'' with $\mathrm{T}_{v''} = \mathrm{T}_v \cup \mathrm{T}_{v'}$ and $\mathrm{F}_{v''} = \mathrm{F}_v \cup \mathrm{F}_{v'}$ by $v \sqcup v'$.*

Compatibility of two interpretations v and v' means that they do not assign conflicting Boolean truth values to a statement. As a consequence, if v and v' are compatible then they have a common successor with respect to the $\leq_i$-relation. In particular, $v \sqcup v'$ is a well-defined interpretation that is such a successor. We can make the following observation.

Proposition 3 *Let $V \subseteq \mathcal{V}$ be a set of interpretations and $v', v'' \in V$ such that v' and v'' are compatible. Then, $v' \sqcup v'' \in V^{\mathrm{ind}}_{\mathrm{adm}}$.*

Proof. Consider some $s \in \mathrm{TF}_v$ for $v = v' \sqcup v''$ and some $v_2 \in [v]_2$. We have $v'(s) = v(s)$ or $v''(s) = v(s)$. Without loss of generality assume $v'(s) = v(s)$. The conjecture holds as $v' \leq_i v_2$. □

We call a set $V \subseteq \mathcal{V}$ of interpretations *upward-closed* if $v' \sqcup v'' \in V$ for all $v', v'' \in V$ that are compatible. The intuition behind the naming is that the property has an upward effect with respect to the information ordering in the sense that v' and v'' imply the presence of $v' \sqcup v''$ which lies higher than v and v' in the corresponding lattice (as in Figure 1).

The next result follows from Theorem 1 and Proposition 3.

Proposition 4 *For every ADF D, $\mathrm{adm}(D)$ is upward-closed.*

If we assume non-emptyness and upward-closedness for the given set of interpretations—properties that are easy to check, we only need to consider interpretations v' with $v \leq_i v'$ to determine whether an interpretation v is adm-induced.

Lemma 1 *Let $V \subseteq \mathcal{V}$ be a non-empty upward-closed set of interpretations and $v \in \mathcal{V}$. Then, $v \in V^{\mathrm{ind}}_{\mathrm{adm}}$ iff for every $v_2 \in [v]_2$ there is an interpretation $v' \in V$ with $v \leq_i v' \leq_i v_2$.*

Proof. ($\Rightarrow$) Consider some $v_2 \in [v]_2$ and the set $V' = \{v' \in V \mid v' \leq_i v_2, s \in \mathrm{TF}_v, v(s) = v'(s)\}$. As elements of V' are mutually compatible, there is an interpretation v'_2 such that $\mathrm{T}_{v'_2} = \bigcup_{v' \in V'} \mathrm{T}_{v'}$ and $\mathrm{F}_{v'_2} = \bigcup_{v' \in V'} \mathrm{F}_{v'}$. Note that v'_2 is the result of iteratively joining the elements of V' using the $\sqcup$-operator. Hence, as V is upward-closed and $V' \subseteq V$ we have by Proposition 3 that $v'_2 \in V$. Clearly, $v'_2 \leq_i v_2$. Consider some $s \in \mathrm{TF}_v$. As $v \in V^{\mathrm{ind}}_{\mathrm{adm}}$, there must be some $v' \in V$ with $v' \leq_i v_2$ and $v(s) = v'(s)$. From that we get that $v' \in V'$ and, consequently, $v'_2(s) = v(s)$. It follows that v'_2 is the interpretation we are looking for because also $v'_2 \leq_i v_2$ holds. ($\Leftarrow$) The conjecture follows immediately from Definition 2. □

In this light, we can reformulate the realizability result for admissible semantics as follows.

Theorem 2 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. There is an ADF D such that $\mathrm{adm}(D) = V$ iff*

(i) V is non-empty,

(ii) upward-closed, and

(iii) for every $v \in \mathcal{V}$ such that for every $v_2 \in [v]_2$ there is some $v' \in V$ with $v \leq_i v' \leq_i v_2$ we have $v \in V$.

The first condition makes explicit that every ADF has an admissible interpretation. In particular, it is easy to see that $v_{\mathbf{u}}$ is admissible for every ADF. In contrast to upward-closedness, condition (iii) has an effect that is downward-directed with respect to the information ordering, i.e., the membership of the interpretations v' in V enforces the membership of v that lies on a lower level of the corresponding lattice.

3.2 Complete Semantics

For complete semantics we start with a couple of necessary conditions for sets $\mathrm{com}(D)$, where D is an ADF. The first condition is downward-oriented. We will need the following counterpart of the $\sqcup$-operator.

Definition 7 *Let $v, v' \in \mathcal{V}$ be interpretations. We denote the interpretation v'' with $\mathrm{T}_{v''} = \mathrm{T}_v \cap \mathrm{T}_{v'}$ and $\mathrm{F}_{v''} = \mathrm{F}_v \cap \mathrm{F}_{v'}$ by $v \sqcap v'$.*

Proposition 5 *Let D be an ADF and $v', v'' \in \mathrm{com}(D)$. Then, there is some $v \in \mathrm{com}(D)$ with $v \leq_i v'$ and $v \leq_i v''$.*

Note that $v \leq_i v'$ and $v \leq_i v''$ means that $v \leq_i v' \sqcap v''$.

Example 1 *Assume $S = \{a, b, c\}$ and consider the ADF*

$$D = \{\langle a, (\neg a \wedge \neg b \wedge \neg c) \vee (a \wedge \neg b \wedge \neg c) \vee (a \wedge \neg b \wedge c) \vee (a \wedge b \wedge c)\rangle, \langle b, (\neg a \wedge \neg b \wedge c) \vee (a \wedge b \wedge c)\rangle, \langle c, (a \wedge \neg b \wedge \neg c) \vee (a \wedge b \wedge c)\rangle\}.$$

D has three complete interpretations, $v_{\mathbf{u}}$, abc, and $a\bar{b}$. Regarding Proposition 5, $v_{\mathbf{u}}$ is the interpretation with $v_{\mathbf{u}} \leq_i abc \sqcap a\bar{b}$ that has to be complete because abc and $a\bar{b}$ are. However, the interpretation $abc \sqcap a\bar{b} = a$, is not complete.

The example shows that, unlike conditions (ii) and (iii) of Theorem 2 that fully determine the interpretations implied to be admissible, in Proposition 5, v' and v'' being complete implies only that one from a class of interpretations is complete.

The next proposition can be considered as an upward-directed counterpart of Proposition 5.

Proposition 6 *Let D be an ADF and $v', v'' \in \mathrm{com}(D)$ compatible interpretations. Then, there is some $v \in \mathrm{com}(D)$ with $v' \sqcup v'' \leq_i v$.*

The difference to upward-closedness is that, similar as in Proposition 5, the implied complete interpretation v is not uniquely determined by v' and v''.

The next three results, however, allow to infer concrete interpretations to be complete under special circumstances. If we consider two complete interpretation that share a direct successor with respect to the information ordering, then the latter is also complete.

Proposition 7 *Let D be an ADF and $v', v'' \in \mathrm{com}(D)$ such that $\mathrm{T}_{v'} = \mathrm{T}_{v''} \cup \{s\}$ and $\mathrm{F}_{v''} = \mathrm{F}_{v'} \cup \{s\}$ for some statement $s \in S$. Then, $v' \sqcap v'' \in \mathrm{com}(D)$.*

Intuitively, in contrast to the general case (cf. Proposition 5), for the common direct successor of v' and v'' there is no further statement whose truth value under one of the two-valued extensions of $v' \sqcap v''$ can spoil completeness. A similar argument justifies the following result.

Proposition 8 *Let D be an ADF, $v'_2, v''_2 \in \mathcal{V}_2$ such that $\mathrm{T}_{v'} = \mathrm{T}_{v''} \cup \{s\}$ or $\mathrm{T}_{v''} = \mathrm{T}_{v'} \cup \{s\}$ for some statement $s \in S$, and $v'_2, v'_2 \sqcap v''_2 \in \mathrm{com}(D)$. Then, $v''_2 \in \mathrm{com}(D)$.*

The next necessary condition is upward-directed and infers a two-valued interpretation to be complete if there are $\leq_i$-predecessors that are complete and jointly assign all statements a Boolean truth value. Note that the result is not an immediate consequence of Proposition 6.

Proposition 9 *Let D be an ADF and $V \subseteq \mathrm{com}(D)$ such that $v_2 \in \mathcal{V}_2$, $\mathrm{T}_{v_2} = \bigcup_{v \in V} \mathrm{T}_v$ and $\mathrm{F}_{v_2} = \bigcup_{v \in V} \mathrm{F}_v$. Then, we have $v_2 \in \mathrm{com}(D)$.*

The following example shows that the necessary conditions expressed in Propositions 5-9 are not sufficient to guarantee realizability under complete semantics.

Example 2 *Assume $S = \{a, b, c\}$ and consider the set $V = \{v_{\mathbf{u}}, \bar{a}, b, \bar{c}, \bar{a}b\bar{c}\}$. V fullfills the properties stated for $\mathrm{com}(D)$ in Propositions 5-9, nevertheless it can be shown (using the notion of a* com-characterization *we introduce below) that there is no ADF that has V as its complete interpretations.*

Intuitively, the complete semantics is difficult to capture because there are many possible reasons why an interpretation is or is not complete. For example, in case an interpretation v is complete, for each statement $s \in \mathrm{U}_v$, there must be one two-valued extension v'_2 of v under which s is true and another v''_2 under which s is false. However, these v'_2 and v''_2 are not necessarily complete and thus we do not know to which truth values they map s. Therefore, we cannot identify v'_2 and v''_2 without knowing the ADF. Conversely, if v is an interpretation that is not complete, we do in general not know which statements are the culprit for that. We can, however, formulate necessary and sufficient conditions for a set V to be realizable that are not based on simple pattern-based checking of the presence or absence of interpretations in V. Completeness of an interpretation in an ADF depends on the truth values that its two-valued extensions assign to the acceptance formulas of each statement. We use the notion of a *com-characterization*, defined next, for providing such truth values (encoded as two-valued interpretation) without reference to an ADF.

Definition 8 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. A function $f : \mathcal{V}_2 \to \mathcal{V}_2$ is a* com-characterization *of V if for each $v \in \mathcal{V}$ we have $v \in V$ exactly when for every $s \in S$*

- *$v(s) \neq u$ implies $f(v_2)(s) = v(s)$ for all $v_2 \in [v]_2$ and*
- *$v(s) = u$ implies $s \in \mathrm{T}_{f(v'_2)}$ and $s \in \mathrm{F}_{f(v''_2)}$ for some $v'_2, v''_2 \in [v]_2$.*

Given a com-characterization for V we can construct an ADF that has V as its set of complete interpretations.

Proposition 10 *Let $V \subseteq \mathcal{V}$ be a set of interpretations, f a com-characterization for V, and D_f the ADF where the acceptance formula for each statement s is given by*

$$\varphi_s = \bigvee_{v_2 \in \mathcal{V}_2, s \in \mathrm{T}_{f(v_2)}} \bigwedge_{s' \in \mathrm{T}_{v_2}} s' \wedge \bigwedge_{s' \in \mathrm{F}_{v_2}} \neg s' \quad .$$

Then, $\mathrm{com}(D_f) = V$.

Proof. First, observe that for every two-valued interpretation v_2 and every $s \in S$ we have $f(v_2)(s) = v_2(\varphi_s)$.

($\subseteq$) Let $v \in \mathrm{com}(D_f)$ be an interpretation and $s \in S$ a statement. Consider the case that $s \in \mathrm{U}_v$. Since $v = \Gamma_{D_f}(v)$, by definition of Γ_{D_f}, there must be two two-valued interpretations $v_2, v'_2 \in [v]_2$ such that $v_2(\varphi_s) = \mathbf{t}$ and $v'_2(\varphi_s) = \mathbf{f}$. By our initial observation, we get $s \in \mathrm{T}_{f(v_2)}$ and $s \in \mathrm{F}_{f(v'_2)}$. Now, consider the case that $s \in \mathrm{TF}_v$. Let v_2 be a two-valued interpretation with $v_2 \in [v]_2$. Since $v = \Gamma_{D_f}(v)$ we have $v(s) = v_2(\varphi_s)$. Therefore, by our observation it must also hold that $f(v_2)(s) = v(s)$. Thus we have $v \in V$ as for every $s \in S$, v meets the respective condition of Definition 8.

($\supseteq$) Consider an interpretation v such that $v \notin \mathrm{com}(D_f)$. We will show that $v \notin V$. As $v \notin \mathrm{com}(D_f)$ we have $v \neq \Gamma_{D_f}(v)$. Let $s \in S$ be a statement such that $v(s) \neq \Gamma_{D_f}(v)(s)$. Consider the case that $s \in \mathrm{TF}_v$. As $v \neq \Gamma_{D_f}(v)$, there must be some $v_2 \in [v]_2$ with $v_2(\varphi_s) \neq v(s)$ and consequently $f(v_2)(s) \neq v(s)$. Thus, by Definition 8 we have $v \notin V$. Now consider the case that $s \in \mathrm{U}_{v(s)}$. As then $\Gamma_{D_f}(v)(s) \neq \mathbf{u}$ we have that for some $\mathbf{x} \in \{\mathbf{t}, \mathbf{f}\}$ we have $v_2(\varphi_s) = \mathbf{x}$ for each $v_2 \in [v]_2$. But then it must also hold that $f(v_2)(s) = \mathbf{x}$ for each $v_2 \in [v]_2$. Again, by Definition 8 we have $v \notin V$. □

Note that, unlike the canonical ADF for a set V that we defined for admissible semantics, the ADF D_f is not unique for V but depends on a concrete f.

Example 3 *Let D be the ADF from Example 1. A com-characterization of $\mathrm{com}(D)$ is given by f:*

v_2	abc	$a\bar{b}c$	$a\bar{b}\bar{c}$	$\bar{a}bc$	$\bar{a}b\bar{c}$	$\bar{a}\bar{b}c$	$\bar{a}\bar{b}\bar{c}$
$f(v_2)$	abc	$a\bar{b}\bar{c}$	$a\bar{b}c$	$\bar{a}\bar{b}\bar{c}$	$\bar{a}\bar{b}\bar{c}$	$\bar{a}b\bar{c}$	$a\bar{b}\bar{c}$

Note that D corresponds to D_f as obtained in Proposition 10.

The existence of a com-characterization is not only a sufficient but also a necessary condition for realizability.

Theorem 3 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. There is an ADF D such that $\mathrm{com}(D) = V$ iff there is a com-characterization for V.*

Proof. *(Sketch)* The 'if'-direction follows from Proposition 10. It remains to show that for any ADF D there is a com-characterization for $\mathrm{com}(D)$. In particular, we define the function f_D as $f_D(v_2)(s) = v_2(\varphi_s)$ for every $v_2 \in \mathcal{V}_2$ and $s \in S$ with $\langle s, \varphi_s \rangle \in D$. The remainder of the proof shows $\mathrm{com}(D) = V$ along the lines as $\mathrm{com}(D_f) = V$ was shown in the proof of Proposition 10, but now the argument of the initial observation of the former proof (that is based on the structure of D_f) is provided by the definition of f_D. □

One could ask whether checking the existence of a com-characterization for checking whether V is realizable has advantages over directly checking the definition of realizability itself, i.e., enumerating all ADFs and checking $\mathrm{com}(D) = V$. In order to answer this question, we next give a complexity result. For the theorem we assume that the set S of all statements is not predefined.

Theorem 4 *Given a finite, non-empty set S as the set of all statements and a set $V' \subseteq \mathcal{V}$ of interpretations, the problem of deciding whether there is an ADF D such that $\mathrm{com}(D) = V'$ is in* NEXPTIME.

The proof of Theorem 4 is based on guessing a candidate for a com-characterization and checking whether it fulfills the conditions of Definition 8. The result shows that we do much better than in a naïve approach where first an ADF D is guessed and then $V = \mathrm{com}(D)$ is tested by checking $v \in \mathrm{com}(D)$ for every interpretation. Remember that there are 3^n many interpretations where n is the cardinality of S and note that checking whether a given interpretation is complete is hard for the complexity class DP (as shown by Strass and Wallner [2014]) which necessitates an NP-oracle call for each interpretation.

3.3 Grounded Semantics

As shown by Brewka *et al.* [2013], every ADF has a single grounded interpretation, in particular that is the minimal complete interpretation with respect to the information ordering. The existence of a unique minimal complete interpretation is also reflected by Proposition 5.

Lemma 2 *For every $v \in \mathcal{V}$, f_v^s is a com-characterization for $\{v\}$, where*

$$f_v^s(v_2)(s) = \begin{cases} v(s) & \text{if } s \in \mathrm{TF}_v \text{ and } v \leq_i v_2 \\ \neg v_2(s) & \text{else} \end{cases}$$

The lemma implies that every singleton set of statements can be realized under complete semantics. For each ADF realizing such a set, grounded and complete semantics coincide. Therefore, also constructing such an ADF is similar as under complete semantics. We get the following realizability result for grounded semantics.

Theorem 5 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. There is an ADF D such that $\mathrm{grd}(D) = V$ iff V has cardinality 1.*

Note that singleton sets of interpretations can be realized under all four semantics we consider.

3.4 Preferred Semantics

While the grounded interpretation is the minimal complete interpretation of an ADF, the preferred interpretations are those complete interpretations that are maximal with respect to the information order. At the same time they are the maximal admissible interpretations [Brewka *et al.*, 2013]. From this we get that the set of preferred interpretations must be *incomparable* as defined next, since the set of admissible interpretations is upward-closed.

Definition 9 *A set $V \subseteq \mathcal{V}$ of interpretations is* incomparable *if all two $v', v'' \in V$ with $v' \neq v''$ are not compatible.*

For selecting the $\leq_i$-maximal interpretations from a set $V \subseteq \mathcal{V}$ of interpretations, we introduce the notation $V^\top = \{v \in V \mid v \leq_i v' \text{ for } v' \in V \text{ implies } v = v'\}$.

The realizability result for preferred semantics can then be given as follows.

Theorem 6 *Let $V \subseteq \mathcal{V}$ be a set of interpretations. There is an ADF D such that $\mathrm{pre}(D) = V$ iff V is non-empty, incomparable, and $V = V^\top$.*

For constructing an ADF for a given V that satisfies the conditions of Theorem 6, the idea is to create an adm-closed set V' with $V = V'^\top$. Then, the canonical ADF $\mathcal{CAN}(V')$ for V' has V as its set of preferred interpretations. In order to obtain V' it suffices to start with V and iteratively add adm-induced interpretations. In this process, whenever, the intermediate set V'' of interpretations is upward-closed, an arbitrary adm-induced interpretation can be added without changing the $\leq_i$-maximal interpretations by Lemma 1. Whenever V'' is not upward-closed, there is some $v', v'' \in V''$ with $v' \sqcup v'' \notin V''$. Then, one can safely add $v' \sqcup v''$ which is admissible induced by V'' and can be shown not to influence the $\leq_i$-maximal interpretations either.

4 Semantical Consideration on Joining ADFs

Having characterizations for the realizability of a semantics is a starting point for studying the expressiveness of a formalism. For example, Strass [2015] uses realizability results to compare the expressiveness of ADFs, AFs, and logic programs under two-valued semantics. An application in answer-set programming is to decide whether a given program can be replaced by another from a syntactically simpler class and to construct the latter if possible [Eiter *et al.*, 2013]. Our results are targeted towards similar applications, e.g., deciding whether a given ADF can be replaced by a bipolar ADF [Brewka and Woltran, 2010] with the same semantics or an AF under the corresponding labeling semantics [Caminada and Gabbay, 2009].

In this section we deal with another interesting problem that can be addressed using the new characterizations and is related to composing ADFs. In particular, we deal with the question whether for a given semantics σ, there can be a join operator $\otimes : \mathcal{D}^2 \to \mathcal{D}$ on ADFs such that for every two ADFs D_1 and D_2 we have $\sigma(D_1 \otimes D_2) = \sigma(D_1) \cap \sigma(D_2)$. One could argue that this property, i.e., that the result of a join accepts exactly those interpretations that are accepted by both operands, is a reasonable requirement for a join operation in specific settings. As an example, consider cases where the two operand ADFs represent two separate discussions and the join should give a cautious summary on common conclusions.[1] The following result shows that there is such a desired $\otimes$-operator for admissible semantics.

Theorem 7 *Let D_1 and D_2 be ADFs. Then, there is an ADF D such that $\mathrm{adm}(D) = \mathrm{adm}(D_1) \cap \mathrm{adm}(D_2)$.*

The proof is based on showing that the intersection of two adm-closed sets is adm-closed.

For complete semantics, on the other hand, we get a negative result. That is, given two ADFs D_1 and D_2, there might be no third ADF D such that $\mathrm{com}(D) = \mathrm{com}(D_1) \cap \mathrm{com}(D_2)$. The following ADFs with three statements provide

[1]Note that different operators for composing ADFs are appropriate for different purposes as argued by Gaggl and Strass [2014].

a counterexample for the general case:

$$\begin{aligned}D_1 = \{&\langle a, a \land \neg b \land \neg c\rangle, \langle b, \neg a \land b \land \neg c\rangle,\\ &\langle c, (a \land \neg b \land \neg c) \lor (\neg a \land b \land \neg c)\rangle\}\\ D_2 = \{&\langle a, a \land \neg b \land c\rangle, \langle b, \neg a \land b \land c\rangle,\\ &\langle c, \neg a \land \neg b \land \neg c\rangle\}\end{aligned}$$

We have that $\mathrm{com}(D_1) = \{v_{\mathbf{u}}, \bar{a}, \bar{b}, \bar{a}\bar{b}\bar{c}\}$ and $\mathrm{com}(D_2) = \{v_{\mathbf{u}}, \bar{a}, \bar{b}, \bar{a}\bar{b}\}$. However, there is no ADF D with $\mathrm{com}(D) = \{v_{\mathbf{u}}, \bar{a}, \bar{b}\}$ because by Proposition 6, a common $\leq_i$-successor of $\bar{a}$ and $\bar{b}$ would be missing.

As a side node, if S is restricted to at most two statements, there exists an ADF D with $\mathrm{com}(D) = \mathrm{com}(D_1) \cap \mathrm{com}(D_2)$ whenever $\mathrm{com}(D_1) \cap \mathrm{com}(D_2)$ is non-empty.

For grounded semantics, given Theorem 5, it is easy to see that there is only an ADF that realizes $\mathrm{grd}(D_1) \cap \mathrm{grd}(D_2)$ if $\mathrm{grd}(D_1)$ and $\mathrm{grd}(D_2)$ coincide.

Finally, for preferred semantics, a join fulfilling the desired requirement is possible under the condition that the intersection of preferred interpretations is non-empty.

Theorem 8 *Let D_1 and D_2 be ADFs. Then, there is an ADF D such that $\mathrm{pre}(D) = \mathrm{pre}(D_1) \cap \mathrm{pre}(D_2)$ iff $\mathrm{pre}(D_1) \cap \mathrm{pre}(D_2) \neq \emptyset$.*

5 Discussion and Conclusion

In this paper we addressed the question whether for a given set of three-valued interpretations, there is an ADF that realizes it. Our results lay the ground for investigating the expressiveness of ADFs under the four semantics we considered. Interestingly, works on realizability for AFs [Dunne *et al.*, 2014; Baumann *et al.*, 2014; Dyrkolbotn, 2014] do not cover complete semantics that turned out to be the hardest to characterize in our setting. For AFs, two variants of realizability have been considered: strict realizability [Dunne *et al.*, 2014] that corresponds to the problems we studied, where the arguments that may be used in the realizing AF are pre-determined, and weaker forms [Baumann *et al.*, 2014; Dyrkolbotn, 2014] where auxiliary arguments may be used. Considering auxiliary statements is also an interesting topic for further work on realizability for ADFs. A natural next step is investigating realizability for important subclasses of ADFs such as bipolar ADFs or AFs under labeling semantics. Together with our current results this would allow us to determine under what circumstances ADFs can be replaced by syntactically simpler ones.

Acknowledgments

The author wishes to thank Stefan Woltran, Thomas Linsbichler, Hannes Strass, and Stefan Ellmauthaler for informative discussions and useful pointers to related work.

References

[Baroni *et al.*, 2011] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.

[Baumann *et al.*, 2014] Ringo Baumann, Wolfgang Dvorák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. In Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan, editors, *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 69–74. IOS Press, 2014.

[Brewka and Woltran, 2010] Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In Fangzhen Lin, Ulrike Sattler, and Miroslaw Truszczynski, editors, *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR'10)*. AAAI Press, 2010.

[Brewka *et al.*, 2013] Gerhard Brewka, Hannes Strass, Stefan Ellmauthaler, Johannes Peter Wallner, and Stefan Woltran. Abstract dialectical frameworks revisited. In Francesca Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*. IJCAI/AAAI Press, 2013.

[Caminada and Gabbay, 2009] Martin W. A. Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2-3):109–145, 2009.

[Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.

[Dunne *et al.*, 2014] Paul E. Dunne, Wolfgang Dvorák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*. AAAI Press, 2014.

[Dyrkolbotn, 2014] Sjur Kristoffer Dyrkolbotn. How to argue for anything: Enforcing arbitrary sets of labellings using afs. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*. AAAI Press, 2014.

[Eiter *et al.*, 2013] Thomas Eiter, Michael Fink, Jörg Pührer, Hans Tompits, and Stefan Woltran. Model-based recasting in answer-set programming. *Journal of Applied Non-Classical Logics*, 23(1-2):75–104, 2013.

[Gaggl and Strass, 2014] Sarah Alice Gaggl and Hannes Strass. Decomposing abstract dialectical frameworks. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA'14)*, volume 266 of *Frontiers in Artificial Intelligence and Applications*, pages 281–292. IOS Press, 2014.

[Gogic *et al.*, 1995] Goran Gogic, Henry A. Kautz, Christos H. Papadimitriou, and Bart Selman. The comparative linguistics of knowledge representation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 862–869. Morgan Kaufmann, 1995.

[Strass and Wallner, 2014] Hannes Strass and Johannes Peter Wallner. Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*. AAAI Press, 2014.

[Strass, 2015] Hannes Strass. The relative expressiveness of abstract argumentation and logic programming. In Blai Bonet and Sven Koenig, editors, *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI'15)*, pages 1625–1631. AAAI Press, 2015.