

INSTITUT FÜR LOGIC AND COMPUTATION
FORSCHUNGSBEREICH FÜR DATENBANKEN UND ARTIFICIAL INTELLIGENCE

**Weighted Abstract
Dialectical Frameworks:
Extended and Revised Report**

DBAI-TR-2018-110

**Gerhard Brewka Jörg Pührer Hannes Strass
Johannes P. Wallner Stefan Woltran**

Institut für Logic and Computation
Forschungsbereich für
Datenbanken und
Artificial Intelligence
Technische Universität Wien
Favoritenstr. 9
A-1040 Vienna, Austria
Tel: +43-1-58801-18403
Fax: +43-1-58801-918403
sek@dbai.tuwien.ac.at
www.dbai.tuwien.ac.at

DBAI TECHNICAL REPORT
2018



TECHNISCHE
UNIVERSITÄT
WIEN

Weighted Abstract Dialectical Frameworks: Extended and Revised Report

Gerhard Brewka, Hannes Strass¹
Jörg Pührer, Johannes P. Wallner, Stefan Woltran²

Abstract. Abstract Dialectical Frameworks (ADFs) generalize Dung’s argumentation frameworks allowing various relationships among arguments to be expressed in a systematic way. We further generalize ADFs so as to accommodate arbitrary acceptance degrees for the arguments. This makes ADFs applicable in domains where both the initial status of arguments and their relationship are only insufficiently specified by Boolean functions. We define all standard ADF semantics for the weighted case, including grounded, preferred and stable semantics. We illustrate our approach using acceptance degrees from the unit interval and show how other valuation structures can be integrated. In each case it is sufficient to specify how the generalized acceptance conditions are represented by formulas, and to specify the information ordering underlying the characteristic ADF operator. We also present complexity results for problems related to weighted ADFs.

¹Leipzig University, Computer Science Institute, Leipzig, Germany.
E-mail: brewka@informatik.uni-leipzig.de, hannes.strass@googlemail.com

²TU Wien, Institute of Logic and Computation, Vienna, Austria.
E-mail: jpuehrer,wallner,woltran@dbai.tuwien.ac.at

This is an extended and corrected version of the paper *Weighted Abstract Dialectical Frameworks* published in the Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018). The conference paper was based on a definition of completions suitable for flat information orderings only. In particular, Prop. 1 in the AAAI paper only holds under the restriction to flat orderings. We thank Bart Bogaerts for pointing this out to us. The new definition of completions used in this report (Def. 3.5) does not suffer from these shortcomings.

1 Introduction

Computational models of argumentation are a highly active area of current research. The field has two main subareas, namely logic-based argumentation and abstract argumentation. The former studies the structure of arguments, how they can be constructed from a given formal knowledge base, and how they logically interact with each other. The latter, in contrast, assumes a given set of abstract arguments together with specific relations among them. The focus is on evaluating the arguments based on their interactions with one another. This evaluation typically uses a specific semantics, thus identifying subsets of the available arguments satisfying intended properties so that the chosen set arguably can be viewed as representing a coherent world view.

In the abstract approach, Dung’s argumentation frameworks (AFs) [23] and their associated semantics are widely used. In a nutshell, an AF is a directed graph with each vertex being an abstract argument and each directed edge corresponding to an attack from one argument to another. These attacks are then resolved using appropriate semantics. The semantics are typically based on two important concepts, namely conflict-freeness and admissibility. The former states that if there is a conflict between two arguments, i.e. one argument attacks the other, then the two cannot be jointly accepted. The latter specifies that every set of accepted arguments must defend itself against attacks. A variety of semantics has been defined, ranging from Dung’s original complete, preferred, stable, and grounded semantics to the more recent ideal and cf2 semantics. The different semantics reflect different intuitions about what “coherent world view” means in this context, see e.g. [7] for an overview.

Despite their popularity, there have been various attempts to generalize AFs as many researchers felt a need to cover additional relevant relationships among arguments (see e.g. the work of [16]). One of the most systematic and flexible outcomes of this research are abstract dialectical frameworks (ADFs) [13, 11, 12]. ADFs allow for arbitrary relationships among arguments. In particular, arguments can not only attack each other, they also may provide support for other arguments and interact in various complex ways. This is achieved by adding explicit acceptance conditions to the arguments which are most naturally expressed in terms of a propositional formula (with atoms referring to parent arguments). This way, it is possible to specify individually for a particular argument, say, under what conditions the available supporting arguments outweigh the counterarguments. Meanwhile various applications of ADFs have been presented, for instance in legal reasoning [1, 2] and text exploration [15]. A mobile argumentation app based on ADF techniques was developed by [38].

The operator-based semantics of ADFs can be traced back to work on approximation fixpoint theory (AFT) [19, 18, 20], an algebraic framework for studying semantics of knowledge representation formalisms. We refer to the work of [39] for a detailed analysis of the relationship between ADFs and AFT. The presentation of our approach in this paper does not assume specific background knowledge in AFT.

The motivation for the work presented here is as follows. The definition of the various ADF semantics is based on an analysis in terms of partial two-valued (or, equivalently, three-valued) interpretations. The output provided by ADFs (and AFs, for that matter) is thus restricted to three options: an argument either is true (accepted) in an intended interpretation, or it is false (rejected),

or its value is unknown. However, many situations in argumentation call for more fine-grained distinctions (see, e.g. [3] for an application of weighted argumentation in the Twitter domain). For instance, it is sometimes natural to assume numerical acceptance degrees, say, taken from the unit interval, and to explore the effect of these degrees on other arguments. The availability of such acceptance degrees allows for new, interesting types of queries to be asked. For instance, under a given semantics (stable, preferred, complete, ...), we may want to know whether the value of a particular argument s is above/below a certain threshold in some or all interpretations of the required type. It also may be useful to be able to distinguish among a finite number of acceptance degrees, say strong accept, weak accept, neutral, weak reject and strong reject. Or it may even be useful to operate on intervals of acceptance degrees.

The goal of this paper is to show how the ADF approach (and thus AFs) can accommodate such acceptance degrees. To put it differently, we aim to bridge two rich research areas, multi-valued logics on the one hand and computational models of argumentation on the other.

We start with the necessary ADF background in Section 2. We then introduce our general framework for weighted ADFs in Section 3. Section 4 focuses on ADFs with acceptance degrees in the unit interval. Section 5 applies the same idea to three other valuation structures. Complexity results for various problems related to weighted ADFs are presented in Section 6. Section 7 discusses related work and concludes.

2 Background

An ADF is a directed node-labelled graph (S, L, C) whose nodes represent statements. The links in L represent dependencies: the status of a node s only depends on the status of its parents (denoted $par(s)$), that is, the nodes with a direct link to s . In addition, each node s is labelled by an associated acceptance condition C_s specifying the conditions under which s is acceptable, whence $C = \{C_s\}_{s \in S}$. Formally, the acceptance condition C_s of node s with parents $par(s)$ is a function $C_s : (par(s) \rightarrow \{\mathbf{t}, \mathbf{f}\}) \rightarrow \{\mathbf{t}, \mathbf{f}\}$. It is convenient to represent the acceptance conditions as a collection $\Phi = \{\varphi_s\}_{s \in S}$ of propositional formulas (using atoms from $par(s)$ and connectives \wedge, \vee, \neg). Then, for any interpretation $w : par(s) \rightarrow \{\mathbf{t}, \mathbf{f}\}$, we have $C_s(w) = w(\varphi_s)$, that is, the acceptance condition C_s evaluates w just like w evaluates φ_s . This leads to the logical representation of ADFs we will frequently use, where an ADF is a pair (S, Φ) with the set of links L implicitly given as $(a, b) \in L$ iff a appears in φ_b .

Semantics assign to ADFs a collection of partial two-valued interpretations, i.e. mappings of the statements to values $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$ where \mathbf{u} indicates that the value is undefined. Mathematically such interpretations are equivalent to 3-valued interpretations, but for the purposes of this paper it is beneficial to view them (interchangeably) also as partial interpretations. The three values are partially ordered by \leq_i according to their information content: \leq_i is the \subseteq -least partial order containing $\mathbf{u} \leq_i \mathbf{t}$ and $\mathbf{u} \leq_i \mathbf{f}$. As usual we write $v_1 <_i v_2$ whenever $v_1 \leq_i v_2$ and not $v_2 \leq_i v_1$. The information ordering \leq_i extends in a straightforward way to partial interpretations v_1, v_2 over S in that $v_1 \leq_i v_2$ if and only if $v_1(s) \leq_i v_2(s)$ for all $s \in S$.

A partial interpretation v is total if all statements are mapped to \mathbf{t} or \mathbf{f} . For interpretations v

and w , we say that w extends v iff $v \leq_i w$. We denote by $[v]_2$ the set of all completions of v , i.e. total interpretations that extend v .

For an ADF $D = (S, L, C)$, statement $s \in S$ and a partial interpretation v , the characteristic operator Γ_D is given by

$$\Gamma_D(v)(s) = \begin{cases} \mathbf{t} & \text{if } C_s(w) = \mathbf{t} \text{ for all } w \in [v]_2, \\ \mathbf{f} & \text{if } C_s(w) = \mathbf{f} \text{ for all } w \in [v]_2, \\ \mathbf{u} & \text{otherwise.} \end{cases}$$

That is, the operator returns an interpretation mapping a statement s to \mathbf{t} (resp. \mathbf{f}) if and only if all two-valued interpretations extending v evaluate φ_s to true (resp. false). Intuitively, Γ_D checks which truth values can be justified based on the information in v and the acceptance conditions.

Given an ADF $D = (S, L, C)$, a partial interpretation v is *grounded* with respect to D if it is the least fixpoint of Γ_D ; it is *admissible* with respect to D if $v \leq_i \Gamma_D(v)$; it is *complete* with respect to D if $v = \Gamma_D(v)$; it is a *model* of D if it is complete and total; it is *preferred* with respect to D if v is maximally admissible with respect to \leq_i . As shown in [11] these semantics generalize the corresponding notions defined for AFs. For $\sigma \in \{\text{adm}, \text{com}, \text{prf}\}$, $\sigma(D)$ denotes the set of all admissible (resp. complete, preferred) interpretations with respect to D .

Example 2.1 Given ADF D over $\{a, b\}$ with $\varphi_a = a \vee \neg b$, $\varphi_b = \neg a$, and interpretations

$$\begin{aligned} v_1 &= \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}, \\ v_2 &= \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}\}, \\ v_3 &= \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}\}, \\ v_4 &= \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}\}. \end{aligned}$$

We get $\text{adm}(D) = \{v_1, v_2, v_3, v_4\}$, $\text{com}(D) = \{v_1, v_3, v_4\}$ (note that $\Gamma_D(v_2) = v_3$, and thus $v_2 \notin \text{com}(D)$), and $\text{prf}(D) = \{v_3, v_4\}$. \diamond

3 The General Framework

In this section we introduce weighted ADFs (wADFs). More precisely, we introduce a general framework which allows us to define wADFs over a chosen set V of values (acceptance degrees) based on an information ordering \leq_i on $V \cup \{\mathbf{u}\}$.¹

Definition 3.1 A *weighted ADF (wADF)* over V is a tuple $D = (S, L, C, V, \leq_i)$, where

- S is a set (of nodes, statements, arguments; anything one might accept or not),
- $L \subseteq S \times S$ is a set of links,

¹Slightly abusing notation we write \leq_i for both the specific ADF ordering ($\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}, \leq_i$) and the generic ordering used here.

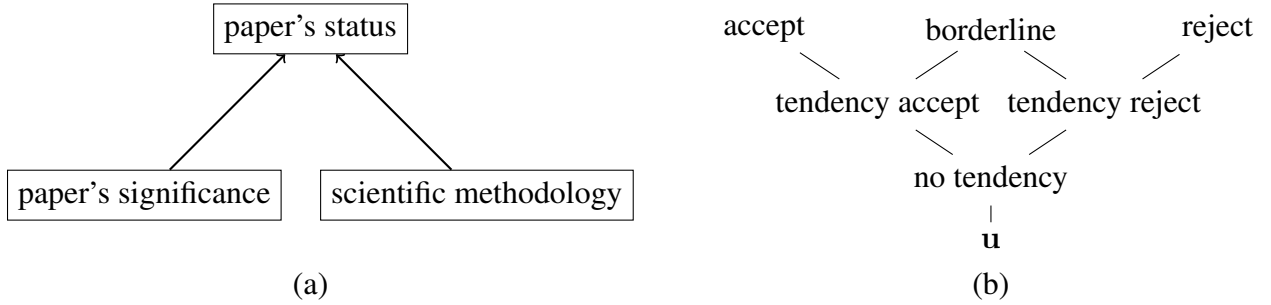


Figure 1: Example wADF (a) and example $V_{\mathbf{u}}$ (b)

- V is a set of truth values with $\mathbf{u} \notin V$,
- $C = \{C_s\}_{s \in S}$ is a collection of acceptance conditions over V , that is, functions $C_s : (par(s) \rightarrow V) \rightarrow V$,
- $(V_{\mathbf{u}}, \leq_i)$ – where $V_{\mathbf{u}} = V \cup \{\mathbf{u}\}$ – forms a complete partial order with least element \mathbf{u} .

As for standard ADFs, the special value \mathbf{u} represents an undefined truth value. As usual, $(V_{\mathbf{u}}, \leq_i)$ forms a complete partial order (CPO) iff: (1) it has a least element, here $\mathbf{u} \in V_{\mathbf{u}}$, (2) each non-empty subset $X \subseteq V_{\mathbf{u}}$ has a greatest lower bound $\prod_i X \in V_{\mathbf{u}}$, and (3) each ascending chain $x_1 \leq_i x_2 \leq_i \dots$ over $V_{\mathbf{u}}$ has a least upper bound $\bigsqcup_i X \in V_{\mathbf{u}}$.

ADFs are a special case of wADF with $V = \{\mathbf{t}, \mathbf{f}\}$ and the information ordering as defined in the background section. We provide a formal result in Theorem 3.19 below.

Example 3.2 In Figure 1a a simple wADF with three arguments is shown that are intended to decide the acceptance status of a paper based on that paper's significance and scientific methodology. On the right side of that figure (b), a value ordering $V_{\mathbf{u}}$ is shown, with the intended meaning that \mathbf{u} denotes no knowledge w.r.t. the arguments, tendencies denote a certain leaning towards acceptance or rejection, and the information maximal values denote acceptance, borderline acceptance, or rejection.

As for ADFs, we will use propositional formulas φ_s interpreted over V to specify acceptance conditions. The understanding is that a formula φ_s specifies a function C_s such that for each interpretation $w : par(s) \rightarrow V$, $C_s(w)$ is obtained by considering $w(\varphi_s)$, the evaluation of the formula φ_s under the interpretation w . Unlike in classical propositional logic, there is no single standard way of interpreting formulas in the multi-valued case. Thus the user (specifying the wADF) should state how formulas are to be evaluated under interpretations of atoms by values from V .

Example 3.3 Continuing Example 3.2, let us define acceptance conditions for each argument. Say argument “paper's significance” (shortened to s) shall be set to “accept” and that the paper's “scientific methodology” (shortened to m) shall be set to “borderline” (e.g. because of peer reviewing). This can be expressed simply by stating: $\varphi_s = \text{accept}$, $\varphi_m = \text{borderline}$. The third

argument, shortened to a , shall depend on the status of the two other arguments, namely by taking the most informative value “compatible” (w.r.t. information ordering) to the values given to the other arguments. That is, take the usual greatest lower bound for two values w.r.t. the information ordering shown in Figure 1b. Say we formalize this by $\varphi_a = s \wedge m$, defining the conjunction as the meet.

In case the truth values in V are \leq_i -incomparable, the information ordering on the truth values $V_{\mathbf{u}} = V \cup \{\mathbf{u}\}$ can be defined analogously to the ordering for standard ADFs (where $V = \{\mathbf{t}, \mathbf{f}\}$ with $\mathbf{t} \not\leq_i \mathbf{f}$ and $\mathbf{f} \not\leq_i \mathbf{t}$).

Definition 3.4 Let V be a set of truth values with $\mathbf{u} \notin V$. A relation $\leq_i \subseteq V_{\mathbf{u}} \times V_{\mathbf{u}}$ is flat iff for all $x, y \in V_{\mathbf{u}}$:

$$x \leq_i y \quad \text{iff} \quad x = \mathbf{u} \text{ or } x = y$$

Likewise, a wADF (S, L, C, V, \leq_i) is flat iff \leq_i is flat.

As mentioned above, clearly all standard ADFs are flat. For flat orderings, the greatest lower bound of a subset $X \subseteq V_{\mathbf{u}}$ is obtained thus:

$$\prod_i X = \begin{cases} x & \text{if } X = \{x\} \\ \mathbf{u} & \text{otherwise} \end{cases}$$

We now define the semantics. A semantics σ takes a wADF D over V and produces a collection $\sigma(D)$ of partial interpretations from S to V , that is, functions $v : S \rightarrow V_{\mathbf{u}}$ with $V_{\mathbf{u}} = V \cup \{\mathbf{u}\}$ where \mathbf{u} represents the fact that the value of a certain node is undefined. Given that for standard ADFs the interpretations of interest are partial functions from S to $\{\mathbf{t}, \mathbf{f}\}$ (or, equivalently, functions from S to $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$), this is the obvious generalization we need.² Let $D = (S, L, C, V, \leq_i)$ be a wADF over V . As for standard ADFs, the characteristic operator for D takes a partial interpretation v and produces a new interpretation, $\Gamma_D(v)$. The new partial interpretation collects information from and mediates between all *completions* of v . As in the standard case a completion of v is any total interpretation w that extends v with respect to the information ordering:³

Definition 3.5 Let $D = (S, L, C, V, \leq_i)$ be a wADF over V , $v : S \rightarrow V_{\mathbf{u}}$ a partial interpretation. A total interpretation $w : S \rightarrow V$ is a completion of v if $v \leq_i w$.

The set of all completions of v is denoted by $[v]_c$.

Example 3.6 Continuing Example 3.2, say we have interpretation $v : \{a \mapsto \text{“tendency accept”}, s \mapsto \text{“accept”}, m \mapsto \text{“borderline”}\}$. Then there are three completions in v , $v_1, v_2 \in [v]_c$: $v_1 : \{a \mapsto \text{“accept”}, s \mapsto \text{“accept”}, m \mapsto \text{“borderline”}\}$ and $v_2 : \{a \mapsto \text{“borderline”}, s \mapsto \text{“accept”},$

²This differs from approaches like [4] which consider weight assignments as part of the input and is more in line with research in multi-valued logics.

³The notion of completion in the original AAI paper considered total interpretations obtained by replacing \mathbf{u} with arbitrary values from V only. This turns out to be insufficient for some of our results to hold.

$m \mapsto \text{“borderline”}$ }. The reason is that both arguments s and m already have information maximal values associated with them (any completion assigns thus the same value), but argument a has a non-maximal value, thus any completion w of v assigns to a any value s.t. $v(a) \leq_i w(a)$. In the example, this means that a is then assigned “tendency accept” (same as for v), “accept” (v_1), and “borderline” (v_2).

Formally, the operator is defined as follows: for each $s \in S$, the truth value $\Gamma_D(v)(s)$ is the greatest lower bound with respect to $(V_{\mathbf{u}}, \leq_i)$ (the consensus) of the set $\{C_s(w) \mid w \in [v]_c\}$. With these specifications the rest is entirely analogous to the definitions for standard ADFs.

Definition 3.7 Let $D = (S, L, C, V, \leq_i)$ be a wADF and $v : S \rightarrow V_{\mathbf{u}}$. Applying Γ_D to v yields a new interpretation (the consensus over $[v]_c$) defined as

$$\Gamma_D(v) : S \rightarrow V_{\mathbf{u}} \quad \text{with} \quad s \mapsto \prod_i \{C_s(w) \mid w \in [v]_c\}.$$

As usual, we can now define the semantics via fixpoints.

Definition 3.8 An interpretation v of a wADF $D = (S, L, C, V, \leq_i)$ is

- grounded for D iff $v = \text{lfp}(\Gamma_D)$, i. e., v is the least fixpoint of Γ_D .
Intuition: v collects all the information which is beyond any doubt.
- admissible for D iff $v \leq_i \Gamma_D(v)$.
Intuition: v does not contain unjustifiable information.
- preferred for D iff it is \leq_i -maximal admissible for D .
Intuition: v has maximal information content without giving up admissibility.
- complete for D iff $v = \Gamma_D(v)$.
Intuition: v contains exactly the justifiable information.
- a model of D iff $v(s) \neq \mathbf{u}$ for all $s \in S$ and $\Gamma_D(v) = v$.
Intuition: v contains exactly the information that is justifiable when each statement has a defined truth value.

Again we use $\text{adm}(D)$, $\text{com}(D)$ and $\text{prf}(D)$ to denote the set of all admissible, complete and preferred interpretations for D , respectively. Moreover, $\text{mod}(D)$ gives the set of all models of D .

Example 3.9 Continuing Example 3.2, let us consider some interpretations. As usual, $I_{\mathbf{u}} = \{a \mapsto \mathbf{u}, s \mapsto \mathbf{u}, m \mapsto \mathbf{u}\}$ is admissible. Since $\Gamma_D(w)(s) = \text{“accept”}$ and $\Gamma_D(w)(m) = \text{“borderline”}$, for any interpretation w (due to definition of their acceptance conditions), it holds that it is admissible to assign these two arguments any value lower or equal to these values (and argument a a “tendency accept”, “no tendency”, or \mathbf{u}). There is only one complete interpretation: $v = \{a \mapsto \text{“tendency accept”}, s \mapsto \text{“accept”}, m \mapsto \text{“borderline”}\}$. This implies that v is the (unique in this case) model of this wADF, and also the only preferred interpretation, as well as the grounded interpretation.

Independently of the previous example, we want to emphasize that we have to show existence of the least fixpoint of Γ_D , otherwise the grounded interpretation is not well-defined in general. The simplest way to do this is to show monotonicity of the operator Γ_D .

To this end, we lift the information ordering on $V_{\mathbf{u}}$ point-wise to interpretations over $V_{\mathbf{u}}$. For $v, w : S \rightarrow V_{\mathbf{u}}$, we set

$$v \leq_i w \quad \text{iff} \quad \forall s \in S : v(s) \leq_i w(s)$$

The pair $(\{v : S \rightarrow V_{\mathbf{u}}\}, \leq_i)$ then forms a CPO in which the characteristic operator Γ_D of wADFs is monotone. The least element of this CPO is the interpretation $v_{\mathbf{u}} : s \mapsto \mathbf{u}$ that maps every statement to \mathbf{u} and we will also use $\bigsqcup_i X$ to denote the least upper bound for subset X of $(\{v : S \rightarrow V_{\mathbf{u}}\}, \leq_i)$.

Proposition 3.10 *The operator Γ_D is \leq_i -monotone, that is: for all interpretations $v, w : S \rightarrow V_{\mathbf{u}}$ we have that $v \leq_i w$ implies $\Gamma_D(v) \leq_i \Gamma_D(w)$.*

Proof. Let $v, w : S \rightarrow V_{\mathbf{u}}$ be two interpretations such that $v \leq_i w$. By definition of $[\cdot]_c$, we find that $[w]_c \subseteq [v]_c$. Thus also $\{C_s(u) \mid u \in [w]_c\} \subseteq \{C_s(u) \mid u \in [v]_c\}$ holds for any $s \in S$. It follows that $\bigsqcup_i \{C_s(u) \mid u \in [v]_c\} \leq_i \bigsqcup_i \{C_s(u) \mid u \in [w]_c\}$, that is, $\Gamma_D(v)(s) \leq_i \Gamma_D(w)(s)$ for any $s \in S$. \square

Existence of the least fixpoint of Γ_D then follows via the fixpoint theorem for monotone operators in complete partial orders (see, e.g., [17], Theorem 8.22).

The following result is a generalization of Theorem 25 of [23] and Theorem 1 by [11].

Theorem 3.11 *Let D be a weighted ADF with an information ordering \leq_i .*

1. *Each preferred interpretation for D is complete, but not vice versa.*
2. *The grounded interpretation for D is the \leq_i -least complete interpretation.*
3. *The complete interpretations for D form a complete meet-semilattice with respect to \leq_i .*

Proof. The first item is shown analogous to a previous result in [39, Theorem 3.10]. More concretely, the first statement in the theorem holds for every approximating operator, such as Γ_D . The second item follows directly from definition (the grounded interpretation is a fixed point, as are all the complete interpretations; the grounded interpretation is the least one). The proof of the third item follows the same line of reasoning as the proof of the third item of [11, Theorem 1]. \square

Next, we show that the well-known relationships between Dung semantics carry over to our generalizations.

Theorem 3.12 *Let D be a weighted ADF. It holds that*

$$\text{mod}(D), \text{prf}(D) \subseteq \text{com}(D) \subseteq \text{adm}(D).$$

If D is flat, then additionally $\text{mod}(D) \subseteq \text{prf}(D)$.

Proof. The “non-flat” case follows from definitions and from Theorem 3.11: every preferred interpretation is complete, every complete interpretation is admissible. Every model is a complete interpretation by definition. For the flat case, every value assigned to a statement by a model is, directly, information maximal (due to the flat ordering). Thus, a model is an information maximal complete interpretation, which directly implies that a model is a preferred interpretation under the assumption that the information ordering is flat. \square

The proviso that $D = (S, L, C, V, \leq_i)$ be flat is necessary for the inclusion $mod(D) \subseteq prf(D)$: consider $S = \{a\}$ with $L = \{(a, a)\}$ and C_a given by $w \mapsto w(a)$ (that is, $\varphi_a = a$); now if there are $x, y \in V$ with $x <_i y$, then we find that $v = \{a \mapsto x\}$ is a model that is not preferred.

The flatness property is also crucial in the following result that lets us compute the grounded semantics by iterative application of the characteristic operator.

Proposition 3.13 *Let $D = (S, L, C, V, \leq_i)$ be a wADF such that every ascending chain in $(\{v : S \rightarrow V_{\mathbf{u}}\}, \leq_i)$ is finite. Then, there is some $n \in \mathbb{N}$ such that $\Gamma_D^n(v_{\mathbf{u}})$ is the grounded interpretation of D . Note that Γ_D^i denotes i iterative applications of Γ_D .*

Proof. As every ascending sequence $v_1 <_i v_2 <_i \dots$ of interpretations is finite and since Γ_D is monotone, it follows from a known result ([17], Theorem 8.8(2)) that Γ_D is Scott-continuous. Then, it follows from the Kleene Fixed-Point Theorem that $v = \bigsqcup_i \{\Gamma_D^i(v_{\mathbf{u}}) \mid i \in \mathbb{N}\}$ is the grounded interpretation of D . As the chain $\Gamma_D(v_{\mathbf{u}}) <_i \Gamma_D(\Gamma_D(v_{\mathbf{u}})) <_i \dots$ is finite, we get that $v = \Gamma_D^n(v_{\mathbf{u}})$ for some $n \in \mathbb{N}$. \square

Note that this result applies to important families of wADFs such as all flat wADFs $D = (S, L, C, V, \leq_i)$ where S is finite or all wADFs $D = (S, L, C, V, \leq_i)$ with finite V .

In general, one cannot guarantee that we reach the grounded interpretation by iterative application of the characteristic operator.

Example 3.14 *Consider wADF $D = (\{s\}, \{(s, s)\}, \{C_s\}, [0, 1], \leq)$, where the acceptance condition of the single statement s is defined as*

$$C_s(v) = \begin{cases} v(s) + \frac{0.5-v(s)}{2} & \text{for } v(s) < 0.5, \\ 1 & \text{otherwise.} \end{cases}$$

We then have

$$\begin{aligned} \Gamma_D(v_{\mathbf{u}})(s) &= 0.25, \\ \Gamma_D^2(v_{\mathbf{u}})(s) &= 0.375, \\ \Gamma_D^3(v_{\mathbf{u}})(s) &= 0.4375, \\ \Gamma_D^4(v_{\mathbf{u}})(s) &= 0.46875, \\ &\dots, \end{aligned}$$

that is, for iterative application of Γ_D on $v_{\mathbf{u}}$, the values converge towards 0.5. However, the interpretation $v = \{s \mapsto 0.5\}$ is not the grounded interpretation of D as it is no fixed point of Γ_D . In fact, $\Gamma_D(v)(s) = 1$. Note that every wADF is guaranteed to have a grounded interpretation, in this case it is $v_g = \{s \mapsto 1\}$.

Despite we do not always get the grounded interpretation by iterative application of the characteristic operator, if we do reach a fixed point this way then it is the grounded interpretation.

Proposition 3.15 *Let D be a wADF. If $v = \bigsqcup_i \{\Gamma^n(v_u) \mid n \in \mathbb{N}\}$ is a fixed point of Γ_D^V , then v is the grounded interpretation of D .*

Proof. The proposition follows from a known result ([17], Theorem 8.15(i)). \square

Another result concerns acyclic wADFs, i.e. ADFs (S, L, C, V, \leq_i) where (S, L) forms an acyclic directed graph and generalizes a recent result of [33].

Theorem 3.16 *For any acyclic wADF D with countable S , $com(D) = prf(D) = \{v\}$ with v the grounded interpretation of D .*

Proof. We show that an acyclic wADF possesses exactly one complete interpretation. The result then follows from Theorem 3.11. Given any acyclic wADF $D = (S, L, C, V, \leq_i)$, we can (partly) order the statements according to their “depth” in the wADF, starting from statements without parents. A path from a statement s_1 to a statement s_n , in D , is defined as $p = (s_1, s_2, \dots, s_{n-1}, s_n)$ with each $s_i \in S$ and for all s_i, s_{i+1} ($1 \leq i \leq n-1$) we have $s_i \in par(s_{i+1})$. The length of p (denoted as $|p|$) is n . We define the depth $d(s)$ of a statement $s \in S$ in D as $d(s) = \max\{|p| \mid p = (s', \dots, s) \text{ is a path in } D, par(s') = \emptyset\}$. Towards a contradiction assume $v_1, v_2 \in com(D)$ such that $v_1 \neq v_2$. Let s be a statement with $v_1(s) \neq v_2(s)$ such that $d(s)$ is minimal (i.e., $v_1(s') = v_2(s')$ for all $s' \in S$ with $d(s') < d(s)$).

Consider the case $d(s) = 1$, where s is a leaf. As $par(s) = \emptyset$, $C_s(w) = v$ maps to the same value v for every interpretation w . Therefore, we have $\bigsqcup_i \{C_s(w) \mid w \in [v_1]_c\} = \bigsqcup_i \{C_s(w) \mid w \in [v_2]_c\}$ which is a contradiction as this is the same as $\Gamma_D(v_1)(s) = \Gamma_D(v_2)(s)$ which amounts to $v_1(s) = v_2(s)$. Now assume $d(s) > 1$. For every parent $p \in par(s)$ we have $d(p) < d(s)$ and thus, by choice of s , it holds that $v_1(p) = v_2(p)$. When we restrict the completions of v_1 and v_2 to $par(s)$, being the domain relevant for evaluating C_s , we get two identical sets of functions $[v_1]'_c = [v_2]'_c$, where $[v_1]'_c = \{s \mapsto w(s) \mid s \in par(s), w \in [v_1]_c\}$ and $[v_2]'_c = \{s \mapsto w(s) \mid s \in par(s), w \in [v_2]_c\}$. One can see that $\bigsqcup_i \{C_s(w) \mid w \in [v_1]_c\} = \bigsqcup_i \{C_s(w) \mid w \in [v_1]'_c\}$ and $\bigsqcup_i \{C_s(w) \mid w \in [v_2]_c\} = \bigsqcup_i \{C_s(w) \mid w \in [v_2]'_c\}$. We end up in the contradiction $\bigsqcup_i \{C_s(w) \mid w \in [v_1]_c\} = \bigsqcup_i \{C_s(w) \mid w \in [v_2]_c\}$ like in the previous case. \square

In the rest of this section we show how stable semantics can be generalized to weighted ADFs. The basic idea underlying stable semantics is to treat truth values asymmetrically. For standard ADFs where only **t** and **f** can appear in models, **f** (false) can be assumed to hold (by default), whereas **t** (true) needs to be justified by a derivation. Technically this is achieved by building the reduct of an ADF and then checking whether the grounded interpretation of the reduct coincides with the original model on the nodes which “survive” in the reduct.

Moving from the two-valued to the multi-valued case offers an additional degree of freedom: it is not clear a priori what the assumed, respectively derived truth values are. The stable semantics we introduce here will thus be parameterized by a subset W of the set of values V over which the weighted ADF is defined.

Definition 3.17 Let $D = (S, L, C, V, \leq_i)$ be a wADF. Let $v : S \rightarrow V$ be a model of D (that is, v is total). Let $W \subseteq V$ be the set of assumed truth values. The v, W -reduct of D is the wADF $D_W^v = (S_W^v, L_W^v, C_W^v, V, \leq_i)$ where

- $S_W^v = \{s \in S \mid v(s) \notin W\}$,
- $L_W^v = L \cap (S_W^v \times S_W^v)$,
- $C_W^v = (C'_s)_{s \in S_W^v}$ where C'_s is obtained from C_s by fixing the value of each parent s^* of s in D such that $s^* \notin S_W^v$ to $v(s^*)$.

Note that whenever the acceptance function is represented using propositional formulas, the new acceptance function is simply obtained by replacing atoms not in S_W^v by their v -values. Now stable models can be defined as usual:

Definition 3.18 Let $D = (S, L, C, V, \leq_i)$ be a wADF and let $v : S \rightarrow V$ be a model of D . Let v_g be the grounded interpretation of the v, W -reduct of D . v is a W -stable model of D iff $v(s) = v_g(s)$ for each $s \in S_W^v$.

This clearly generalizes stable semantics for standard ADFs: just let $V = \{\mathbf{f}, \mathbf{t}\}$ and $W = \{\mathbf{f}\}$. We conclude this section by showing the exact relationship between ADFs and wADFs.

Theorem 3.19 Let $F = (S, L, C)$ be an ADF. The wADF associated to F is $D_F = (S, L, C, \{\mathbf{t}, \mathbf{f}\}, \leq_i)$ with \leq_i as defined in the background section. An interpretation v is a model/admissible/complete/preferred/grounded for F iff it is a model/admissible/complete/preferred/grounded for D_F . Moreover, v is stable for F iff it is $\{\mathbf{f}\}$ -stable for D_F .

Proof. It suffices to show that $\Gamma_{D_F}^{\{\mathbf{t}, \mathbf{f}\}} = \Gamma_F$, i.e., that the characteristic operator for the weighted ADF D_F is the same as the characteristic operator for the non-weighted (standard) ADF $F = (S, L, C)$. To see this, note that D_F recovers all ingredients of standard ADFs: \leq_i is a flat ordering, the same as the “pre-defined” one for F on $\{\mathbf{t}, \mathbf{f}\}$, evaluation of acceptance conditions (or formulas if the representation is different) is the same, and truth values are the same, as well. For the stable semantics, note that the same notion of reduct is recovered for $\{\mathbf{f}\}$ -stable semantics. \square


4 Weighted ADFs Over the Unit Interval

In this section we focus on weighted ADFs over the unit interval, that is, wADFs over $V = [0, 1]$.

We first assume a flat information ordering \leq_i . As discussed in Section 3, we will use propositional formulas φ_s to specify acceptance conditions over $[0, 1]$. In the subsequent examples, we employ a formula evaluation that is defined via structural induction as follows: $w(\varphi \wedge \psi) = \min \{w(\varphi), w(\psi)\}$, $w(\varphi \vee \psi) = \max \{w(\varphi), w(\psi)\}$, and $w(\neg\varphi) = 1 - w(\varphi)$. (Clearly, a richer formula syntax and other evaluations known from multi-valued logics are possible, but not the main topic of this paper.) We furthermore allow (representations of) elements of V

to appear as atoms in the propositional formulas φ_s and let $w(a) = a$ for $a \in V$. This enables us to fix acceptance degrees for specific nodes, and to express upper and lower bounds. For instance, the formula $\phi \wedge 0.7$ expresses that the acceptance degree of a node cannot be higher than 0.7, and similarly $\phi \vee 0.7$ expresses that the acceptance degree cannot be below 0.7.

Example 4.1 Consider wADF D over $[0, 1]$ depicted below.



Intuitively, φ_a fixes the value of a to 0.8. φ_b expresses self-attack. φ_c means c is accepted to the extent a and b are, while d is attacked by b . In addition, it is known, for whatever reason, that the value of d must be at least 0.6.

We represent a partial interpretation v as the tuple $(v(a), v(b), v(c), v(d))$. By Proposition 3.13, the grounded interpretation can be obtained by iterating Γ_D on the interpretation $(\mathbf{u}, \mathbf{u}, \mathbf{u}, \mathbf{u})$. We obtain the least fixpoint $v_1 = (0.8, \mathbf{u}, \mathbf{u}, \mathbf{u})$. The (unique) model of D is $v_2 = (0.8, 0.5, 0.5, 0.6)$. This model is W -stable for $W = [0, 0.5]$. To see this, consider the v_2, W -reduct which consists of nodes a and d with reduced acceptance conditions 0.8 and $0.5 \vee 0.6$, respectively. The grounded interpretation assigns to these nodes exactly the values they have in v_2 , namely 0.8 and 0.6. Note that v_2 is not $[0, 0.5]$ -stable. In this case the reduct is identical to D and the grounded interpretation v_1 of D differs from v_2 .

Interpretations v_1 and v_2 are also the only complete ones. An interpretation v is admissible for D if and only if

1. $v(a) = \mathbf{u}$ or $v(a) = 0.8$,
2. $v(b) = \mathbf{u}$ or $v(b) = 0.5$,
3. $v(c) = \mathbf{u}$ (if $v(a) = \mathbf{u}$ or $v(b) = \mathbf{u}$) or $v(c) = 0.5$ (if $v(a) = 0.8$ and $v(b) = 0.5$), and
4. $v(d) = \mathbf{u}$ (if $v(b) = \mathbf{u}$) or $v(d) = 0.6$ (if $v(b) = 0.5$).

The single preferred interpretation for D is v_2 . ◇

The following lemma will be useful to derive results for wADFs on the unit interval (with flat ordering).

Lemma 4.2 Let φ_s be a propositional formula over variables S (without constants). Assume that the formulas are evaluated on values in the unit interval as specified in the beginning of Section 4. Further, let $v : S \rightarrow \{\mathbf{t}, \mathbf{f}\}$ be a total interpretation assigning variables in S to either true or false, and $w : S \rightarrow [0, 1]$ be a total interpretation assigning variables in S to a value in $[0, 1]$ s.t. $v(s) = \mathbf{t}$ implies $w(s) \geq 0.5$ and $v(s) = \mathbf{f}$ implies $w(s) \leq 0.5$. If v satisfies φ_s then $\varphi_s(w) \geq 0.5$ and if v does not satisfy φ_s then $\varphi_s(w) \leq 0.5$.

Proof. We prove this lemma by structural induction. For the base cases, i.e., atoms and negated atoms, the implications immediately follow. For compound formulas, i.e., $A \circ B$ with $\circ \in \{\wedge, \vee\}$ and $\neg A$, if v satisfies $A \vee B$, then either $v(A)$ or $v(B)$ is **t**, and thus $\max\{w(A), w(B)\} \geq 0.5$. Similarly for \wedge . For negated formulas, if v satisfies $\neg A$, then v does not satisfy A . \square

We next explore the relation between wADFs over the unit interval and ADFs. For the proposition we identify the ADF truth values **t** and **f** with 1 and 0, respectively. We have the following result:

Proposition 4.3 *Let $D = (S, L, C, V, \leq_i)$ be a wADF with no constants other than 0 or 1 appearing in any acceptance formula of C , and $D' = (S, L, C, \{0, 1\}, \leq'_i)$ be its classical version (with $\leq'_i = \leq_i \cap (\{0, 1, \mathbf{u}\} \times \{0, 1, \mathbf{u}\})$). Furthermore, let v be a partial interpretation assigning truth values in $\{0, 1, \mathbf{u}\}$ only, and $s \in S$.*

- *If $\Gamma_D(v)(s) \in \{0, 1\}$, then $\Gamma_{D'}(v)(s) \in \{0, 1\}$.*
- *If $\Gamma_{D'}(v)(s) = \mathbf{u}$, then $\Gamma_D(v)(s) = \mathbf{u}$.*

Proof. Suppose the first item does not hold: $\Gamma_D(v)(s) = 1$ and $\Gamma_{D'}(v)(s) \notin \{0, 1\}$ (symmetric for the case of 0). It holds that $\Gamma'_{D'}(v)(s) = \mathbf{u}$, and thus, there exists a completion of v to v' (on $\{\mathbf{t}, \mathbf{f}\}$) s.t. v' does not satisfy φ_s . Construct an interpretation w' s.t. $w'(s) = 1$ iff $v'(s) = \mathbf{t}$ and $w'(s) = 0$ iff $v'(s) = \mathbf{f}$. We have w' is a completion of v (on the unit interval). By Lemma 4.2, it holds that $\varphi_s(w') \leq 0.5$. This is a contradiction to $\Gamma_D(v)(s) = 1$ (all completions on the unit interval have to evaluate to 1 for φ_s).

The second item of the proposition can be shown analogously. Suppose the contrary, then there exist two completions of v (on $\{\mathbf{t}, \mathbf{f}\}$), s.t. one satisfies φ_s while the other does not. By Lemma 4.2, it holds that the corresponding interpretations on the unit interval (identifying, again, 0 with **f** and 1 with **t**), evaluate once to a value at most 0.5 and otherwise to a value at least 0.5. To see that the end result cannot be 0.5, by definition of the connectives, only values from the interpretation (here only $\{0, 1\}$) can be re-produced when evaluating the connectives. Therefore, the characteristic operator $\Gamma_D(v)(s)$ returns **u**. \square

This result cannot be strengthened, in particular $\Gamma_{D'}(v)(s)$ may be 1 or 0, yet $\Gamma_D(v)(s) = \mathbf{u}$, as illustrated as follows:

Example 4.4 *Consider the graph consisting of nodes a, b with acceptance formulas $\varphi_a = a$ and $\varphi_b = a \vee \neg a$. It is easy to see that in the weighted case the grounded interpretation is $\{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}$. In contrast, the standard approach yields the grounded interpretation $\{a \mapsto \mathbf{u}, b \mapsto 1\}$. This is due to the fact that $a \vee \neg a$ is a tautology in two-valued logic, but not when the unit interval and the above specified formula evaluation is used. Here the formula may have any value $x \geq 0.5$. \diamond*

A possible remedy would be to define non-standard interpretations \wedge^*, \vee^* for the connectives \wedge, \vee , for example:

- $x \wedge^* y = 1$ if $x > 0.5$ and $y > 0.5$; $x \wedge^* y = 0$ otherwise
- $x \vee^* y = 1$ if $x > 0.5$ or $y > 0.5$; $x \vee^* y = 0$ otherwise

However, this approach appears to throw out the baby with the bath water: the different behavior of wADFs in such examples stems from the fact that they make more fine-grained distinctions. It is thus not unintended. Note also that there is an easy alternative option to specify tautological acceptance conditions for wADFs: simply replace $a \vee \neg a$ with 1.

So far we have used a flat information ordering with least element \mathbf{u} and different elements in $[0, 1]$ incomparable. Of course, nothing prevents us from choosing a more refined ordering, for instance the ordering \leq'_i given by

$$x \leq'_i y \quad \text{iff} \quad x \leq_i y \text{ or } y < x \leq 0.5 \text{ or } 0.5 \leq x < y$$

That is, 0.5 is immediately above \mathbf{u} , and a value smaller than 0.5 is more informative if it is closer to 0, a value greater than 0.5 is more informative if it is closer to 1. The pair $([0, 1] \cup \{\mathbf{u}\}, \leq'_i)$ again forms a complete partial order; for any non-empty $X \subseteq [0, 1] \cup \{\mathbf{u}\}$, its greatest lower bound is given by

$$\prod'_i X = \begin{cases} \mathbf{u} & \text{if } \mathbf{u} \in X \\ \min X & \text{if } X \subseteq [0.5, 1] \\ \max X & \text{if } X \subseteq [0, 0.5] \\ 0.5 & \text{otherwise.} \end{cases}$$

The CPO property extends to the pointwise extension of \leq'_i to valuations. Thus for any $D = (S, L, C, [0, 1], \leq'_i)$, its characteristic wADF operator Γ_D is well-defined, in particular its least fixpoint $\text{lfp}(\Gamma_D)$ exists and is uniquely determined.

Example 4.5 Consider again Example 4.1, but this time with information ordering \leq'_i . Again we iterate Γ_D on the interpretation $(\mathbf{u}, \mathbf{u}, \mathbf{u}, \mathbf{u})$. With the extended information ordering we obtain the fixpoint $v = (0.8, 0.5, 0.5, 0.6)$ which is guaranteed to be the grounded interpretation by Proposition 3.15. The admissible semantics is given by all interpretations in

$$\{\{a \mapsto \mathbf{v}_a, b \mapsto \mathbf{v}_b, c \mapsto \mathbf{v}_c, d \mapsto \mathbf{v}_d\} \mid \begin{array}{l} \mathbf{v}_a \in \{\mathbf{u}\} \cup [0.5; 0.8], \\ \mathbf{v}_b \in \{\mathbf{u}, 0.5\}, \\ \mathbf{v}_c \in \{\mathbf{u}, 0.5\}, \\ \mathbf{v}_d \in \{\mathbf{u}\} \cup [0.5; 0.6] \end{array}\}.$$

As every complete interpretation is admissible and v is \leq'_i -minimal among all complete interpretations but, in this case, v coincides with the \leq'_i -maximal admissible interpretation, we get that v is the only complete and preferred interpretation as well as the only model. Similar as in Example 4.1, the v, W -reduct for $W = [0, 0.5]$ consists of nodes a and d with reduced acceptance conditions 0.8 and $0.5 \vee 0.6$, respectively. Thus, v is a W -stable model. Unlike in Example 4.1, v is also W' -stable, for $W' = [0, 0.5[$, since it is the grounded interpretation and the ADF coincides with its v, W' -reduct. \diamond

5 Alternative Valuation Structures

In the last section we considered values in $[0, 1]$. Of course, there are many more options which have been studied intensively in the area of multi-valued logics (an excellent overview was given by [28]), e.g.

- $W_m = \{\frac{k}{m-1} \mid 0 \leq k \leq m-1\}$,
- Belnap's 4-valued system with $\{\emptyset, \{\perp\}, \{\top\}, \{\perp, \top\}\}$, or
- values from an interval $[a, b]$ within the unit interval ($0 \leq a \leq b \leq 1$).

Independently from the actual choice, as before, given $D = (S, L, C, V, \leq_i)$, for any chosen set of truth degrees V , an interpretation assigns a value from $V_{\mathbf{u}}$ to nodes in S , and the characteristic operator Γ_D works on partial V -interpretations, that is, on functions of the type $S \rightarrow V_{\mathbf{u}}$ where $\mathbf{u} \notin V$. Recall that acceptance conditions are of the form $C_s : (par(s) \rightarrow V) \rightarrow V$.

For one choice of values, we represent acceptance conditions as propositional formulas, this time interpreted over V . To make ADF techniques work we need to define the evaluation of propositional connectives over V , thus specifying which acceptance condition a formula actually represents, and the information ordering, as done for $[0, 1]$ in Section 4.

The literature on multi-valued logics provides a rich source of alternative valuation structures with different benefits and properties. It also offers a wide range of options regarding different evaluations of propositional formulas (e. g., Gödel, Łukasiewicz, etc.). The only general constraint is that the information ordering \leq_i on $V \cup \{\mathbf{u}\}$ must form a complete partial order (CPO) with least element \mathbf{u} .

In the following we illustrate the use of alternative valuation structures using three different examples.

W₃: For the family of values W_m , let us exemplify $W_3 = \{0, 0.5, 1\}$. We define formula evaluation as before, that is, 0, 0.5, and 1 evaluate to themselves, \wedge to min, \vee to max, and $\neg y$ to $1 - y$. We choose as the information ordering the smallest reflexive relation containing $\mathbf{u} \leq_i x$ for all $x \in W_3$, where the remaining values are incomparable if they are different. That is, we (again) utilize here a flat ordering. This fully specifies wADFs based on W_3 . Note that $0.5 \neq \mathbf{u}$. This makes perfect sense as saying “the acceptance degree is 0.5” is different from saying “the acceptance degree is unknown”.

Belnap: We show how the truth degrees in Belnap's four-valued logic can be used in wADFs. The truth degrees are

$$\mathbf{B} = \{\emptyset, \{\perp\}, \{\top\}, \{\perp, \top\}\}.$$

As to formula evaluation we use the standard definitions for Belnap's logic: conjunction/disjunction are the infimum/supremum under the truth ordering \leq_t ; negation swaps $\{\perp\}$ and $\{\top\}$, and leaves the other two values unchanged. The truth ordering is the reflexive closure of:

$$\{\perp\} \leq_t \emptyset \leq_t \{\top\} \quad \{\perp\} \leq_t \{\perp, \top\} \leq_t \{\top\}$$

The information ordering is the reflexive closure of:

$$\mathbf{u} \leq_i \emptyset \leq_i \{\top\} \leq_i \{\perp, \top\} \quad \mathbf{u} \leq_i \emptyset \leq_i \{\perp\} \leq_i \{\perp, \top\}$$

With these definitions, the operator Γ_D and thus the 4-valued wADF system is fully specified. Note again that $\mathbf{u} \neq \emptyset$; treating them as identical would yield a different system.

Intervals: Our approach can also handle intervals. Let us illustrate this idea using intervals from within the unit interval as truth degrees, that is we consider truth values in

$$INT = \{[a, b] \mid 0 \leq a \leq b \leq 1\}.$$

Formula evaluation can be defined as follows:

$$\begin{aligned} [a, b] \wedge^* [c, d] &= [\min(a, c), \min(b, d)] \\ [a, b] \vee^* [c, d] &= [\max(a, c), \max(b, d)] \\ \neg^*[a, b] &= [1 - b, 1 - a] \end{aligned}$$

The information ordering is the \subseteq -least relation satisfying (1) $\mathbf{u} \leq_i v$ for all $v \in INT$, and (2)

$$[a, b] \leq_i [c, d] \quad \text{whenever} \quad [c, d] \subseteq [a, b],$$

which fully defines the characteristic operator Γ_D .

6 Computational Complexity

In this section we give a preliminary complexity analysis of weighted ADFs. For this we assume that all acceptance conditions are specified via propositional formulas, which, additionally, may contain constants from a pre-specified and fixed set of values V . Further, we assume that evaluating an acceptance condition under an interpretation that assigns no statement to \mathbf{u} can be done in polynomial time, and that comparing two values under the, again pre-specified and fixed, information ordering can, likewise, be computed in polynomial time with respect to the size of the given wADF. That is, both V and \leq_i are fixed ($|V|$ is a constant) and not part of the input for the problems we study in this section.

We first show that, under fixed and finite valuation structures V , complexity upper bounds of wADFs remain the same as for classical ADFs [41] for several central computational tasks. As for ADFs, a cornerstone auxiliary complexity result is the following for checking whether a given interpretation is admissible in a given wADF.

Proposition 6.1 *Verifying whether a given interpretation is admissible in wADFs with fixed and finite valuation structures is in **coNP**.*

Proof. We show the membership result for the complementary problem. A given interpretation v over set of values V is not admissible in a given wADF $D = (S, L, C, V, \leq_i)$ iff

$v \not\leq_i \Gamma_D(v)$ iff $\exists s \in S$ s.t. $v(s) \not\leq_i \Gamma_D(v)(s)$ iff $v(s) \not\leq_i \prod_i \{C_s(w) \mid w \in [v]_c\}$. It holds that $\{C_s(w) \mid w \in [v]_c\} \subseteq V$ and there exists a subset $W \subseteq [v]_c$ of completions of v s.t. $\{C_s(w) \mid w \in [v]_c\} = \{C_s(w) \mid w \in W\}$ and $|W| \leq |V|$ (we only need one completion per value). By definition we have $|V|$ is constant, and therefore $|W|$ is bounded by a constant.

Non-deterministically guess an $s \in S$, $|V|$ -many completions of v , say as set X , and compute $\prod_i \{C_s(w) \mid w \in X\}$. It holds that $\{C_s(w) \mid w \in X\} \subseteq \{C_s(w) \mid w \in [v]_c\}$. Now, compute the greatest lower bound y wrt the completions X , which acts as an approximation of the greatest lower bound of all completions. The greatest lower bound $y = \prod_i \{C_s(w) \mid w \in X\}$ can be computed in polynomial time, since cardinality of X is bounded by a constant, by assumption evaluation $C_s(w)$ can be computed in polynomial time, and one can check for each value in V whether this value is the greatest lower bound of $\{C_s(w) \mid w \in X\}$. It holds that y approximates $\prod_i \{C_s(w) \mid w \in [v]_c\}$ by: $\prod_i \{C_s(w) \mid w \in [v]_c\} \leq_i y$. The last inequality follows since y is the greatest lower bound of a subset of $\{C_s(w) \mid w \in [v]_c\}$, and the information ordering and values form a CPO; thus y cannot be lower than the greatest lower bound of the superset, and not incomparable, which would violate uniqueness of the greatest lower bound in the superset. Having computed y , we check whether $v(s) \leq_i y$. If $v(s) \not\leq_i y$, then $v(s) \not\leq_i \prod_i \{C_s(w) \mid w \in [v]_c\}$ (v is not admissible). On the other hand, if $v(s) \leq_i \prod_i \{C_s(w) \mid w \in [v]_c\}$, then $v(s) \leq_i y$. Finally, the actual $\prod_i \{C_s(w) \mid w \in [v]_c\}$ can always be found by guessing the corresponding completions.

Thus, the problem of checking whether v is not admissible is in NP, and, therefore, the problem of verifying whether v is admissible in D is in coNP. \square

Based on the previous result, checking whether there exists an admissible interpretation for a given wADF that assigns a given value, different to \mathbf{u} , to a given statement, has the same complexity upper bound as the analogous task of credulous acceptance on ADFs. The complexity class Σ_2^p contains all problems solvable via a non-deterministic polynomial-time algorithm that has access to an NP-oracle (which can solve problems in NP in constant time).

Proposition 6.2 *Checking whether there is an admissible interpretation assigning a given statement a given value for wADFs with fixed and finite valuation structures is in Σ_2^p .*

Proof. For given wADF $D = (S, L, C, V, \leq_i)$, $a \in S$, and $x \in V$, non-deterministically guess a partial interpretation v with $v(a) = x$ and check whether v is admissible in D (a problem in coNP due to Proposition 6.1). \square

Checking whether all preferred interpretations for a given wADF assign a given value, different to \mathbf{u} , to a given statement, has the same complexity upper bound as the analogous task of skeptical acceptance on ADFs. The class Π_3^p is the complement class of Σ_3^p (contains the complements of all problems in Σ_3^p), which in turn contains all problems solvable via a non-deterministic polytime algorithm with access to a Σ_2^p -oracle.

Proposition 6.3 *Checking whether all preferred interpretations assign a given statement a given value for wADFs with fixed and finite valuation structures is in Π_3^p .*

Proof. Let $D = (S, L, C, V, \leq_i)$ be a wADF, $a \in S$, and $x \in V$. To check whether a given interpretation v is preferred in D , consider the complementary problem: checking whether v is not preferred in D . Non-deterministically guess v' with $v \leq_i v'$, and, via an NP-oracle, both check whether v' is admissible in D and whether v is admissible in D . If v' is admissible or v is not admissible, we can directly infer that v is not preferred in D . Thus, checking whether v is preferred in D is in Π_2^p .

To check whether all preferred interpretations of D assign s to x , consider again the complementary problem: to check whether there exists a preferred interpretation assigning s to a value different than x . Non-deterministically guess a partial interpretation v with $v(s) \neq x$ and check whether v is preferred in D via a Π_2^p -oracle. Thus, the problem of verifying whether all preferred interpretations of D assign s to x is in Π_3^p . \square

Analogously to ADFs, the same complexity upper bound for existence of stable models can be derived for wADF with fixed and finite valuation structures. The assumed values can be arbitrarily chosen among the fixed V , but are given as input, as well.

Proposition 6.4 *Checking existence of stable models is in Σ_2^p for wADFs with fixed and finite valuation structures.*

Proof. Assume a given wADF $D = (S, L, C, V, \leq_i)$, with V fixed. We show a polynomial non-deterministic algorithm, with access to an NP oracle, that decides whether D has a W -stable model. First, non-deterministically construct an interpretation v . Observe that computing the result of $\Gamma_R^V(v')(s)$ can be done in polynomial time with access to an NP oracle for any given v' : checking whether $\Gamma_R^V(v')(s) \neq a$, $a \in V$, is in NP (non-deterministically guess a completion and evaluate the acceptance condition). To check whether v is a model of D it is sufficient to check whether $\Gamma_D(v) = v$ and whether v assigns no argument to \mathbf{u} . By the observation before, polynomially many NP calls are sufficient.

Next, consider the v, W -reduct R . It is immediate that the arguments and links of this reduct can be constructed in polynomial time. For the acceptance conditions, it is sufficient to restrict these to the changed parent relation (completions are fixed for non-parents), and, thus, acceptance conditions need not be modified. To verify stability of v , it suffices to compute the grounded interpretation of R . To achieve this, iterate over each statement s and compute $\Gamma_R^V(v')(s)$, with v' being the interpretation assigning each statement to \mathbf{u} . Applying this procedure for each statement until a fixed point is reached can, likewise, be done in polynomial time (with access to an NP oracle). Comparing the grounded interpretation of R to v can be done in polynomial time. \square

Complexity lower bounds depend on the fixed V , information ordering, and formula evaluation. Non-weighted ADFs (i. e., $V = \{\mathbf{t}, \mathbf{f}\}$) are an example where, for the corresponding fixed components, the complexity lower bounds match the previously shown upper bounds.

Finally, we show a result for wADFs over the unit interval. As before, we make the same assumptions on the acceptance conditions, but let V be the unit interval and assume a flat information ordering (\mathbf{u} is strictly lower than all other elements, and all other elements are incomparable), and

assume formula evaluation as defined in Section 4. Although ADFs are not a special case, with respect to all semantics, of such wADFs, **coNP**-hardness of verifying whether a given interpretation is admissible follows from a similar argument as for classical (non-weighted) ADFs.

Proposition 6.5 *Verifying whether a given interpretation is admissible in wADFs over the unit interval with flat information ordering is **coNP**-hard.*

Proof. We reduce the problem of checking whether a given propositional formula ψ over variables B is a tautology. W.l.o.g., we assume that ψ does not contain any constants (i.e. does not contain \top or \perp). Construct wADF $D = (B \cup \{t\}, C)$ with

- $\varphi_b = b$ with $b \in B$; and
- $\varphi_t = \psi \wedge x$ with $0 < x < 0.5$ (arbitrary choice).

The wADF D can be constructed in polynomial time. We claim that v is admissible in D with $v(b) = \mathbf{u}$ and $v(t) = x$ iff ψ is a tautology. Assume ψ is a tautology. Then any total *two-valued* interpretation (assigning only true or false) satisfies ψ . Thus, due to Lemma 4.2, for any total interpretation $w : B \cup \{t\} \rightarrow [0, 1]$ we have $w(\psi) \geq 0.5$, and, in turn, $w(\varphi_t) = x$ ($\min\{x, w(\psi)\}$ always equals x). This means v is admissible in D .

Assume ψ is not a tautology. Then there exists a two-valued interpretation w s.t. $w(\psi) = \mathbf{f}$. We show that v' with

$$v'(b) = \begin{cases} 1 & \text{if } w(b) = \mathbf{t} \\ 0 & \text{if } w(b) = \mathbf{f} \end{cases}$$

for $b \in B$ and $v'(t) = x$ is a completion of v , i.e., $v' \in [v]_c$, and $v'(\varphi_t) = 0$. To show that $v' \in [v]_c$, consider that $v(b) \leq_i v'(b)$ for all $b \in B$ and $v'(t) = v(t)$. Since v' assigns either 0 or 1 to any statement in B it holds that $v'(\psi)$ is either 0 or 1 (all logical connectives are defined via minimum or maximum functions, or $1 - y$). Due to Lemma 4.2, $v'(\psi)$ cannot be 1 (since w does not satisfy ψ). Thus, $v'(\varphi_t) = \min\{0, x\} = 0$. This means, there exists a completion of v , namely v' , and $v(t) = x \neq 0 = v'(t)$, and, in turn, that v is not admissible. \square

7 Related Work and Conclusions

In this paper we introduced a framework for defining weighted ADFs, a generalization of ADFs allowing to assign acceptance degrees to arguments. The framework is fully flexible regarding the choice of acceptance degrees and their associated information ordering. We have provided definitions of the main semantics and showed a number of properties together with a preliminary complexity analysis.

There is quite some work on weighted argumentation frameworks, and even a section entitled “Weighted ADFs” in [13]. In many cases weights of some sort are added to the *links* in the argument graph, not to the nodes. For instance, Brewka and Woltran use weights on links to simplify the definition of acceptance conditions, an idea that has later been extended to the GRAPPA

framework [14]. In [25] weights on links are used as an “inconsistency budget”: attacks may be disregarded as long as the sum of the weights of disregarded attacks remains under some threshold.

Here we focus on papers assigning acceptance degrees to argument nodes. One such approach is Gabbay’s equational theory of argumentation frameworks [27]. He allows for values in the unit interval and represents AFs in an equational form, where the equations specify certain constraints value assignments need to satisfy.

There are also probabilistic extensions of AFs, e.g. [24, 35, 30, 31, 32] (for a complexity analysis for probabilistic AFs see [26]), and even of ADFs [37]. The main idea is to generate several standard AFs (resp. ADFs) which represent the possible situations induced by the probabilities. The latter can be assigned to arguments, attacks, and in case of ADFs, to acceptance conditions. The evaluation of frameworks generated this way follows the standard approach, and the results of these evaluations are aggregated accordingly. The behavior of the semantics is thus triggered via all relevant subgraphs. A related approach in a multi-valued setting is [21, 22].

Social AFs [34] extend standard AFs by adding to each argument associated numbers of positive and negative votes. The semantics describe how the votes propagate through the network, yielding a non-linear system of equations. Recently, several properties and semantics for weighted AFs have been proposed in [6, 5]. In those works, weights for arguments are also given from the unit interval, interpreted in the sense that the greater the value the more acceptable the argument. The focus is on the definition of new semantics dedicated for weighted AFs, rather than on generalizing standard semantics. However, the properties proposed in [5] adapted to our setting are of interest and thus are on our agenda for future work.

Finally in [9] acceptance grades of arguments are derived from the structure of the argument tree. The authors attempt to “*provide an abstraction of an argument tree in the form of a single number*”. In a similar vein Grossi and Modgil [29] derive acceptance grades from the structure of the underlying AF, e.g. the number of attacks against which a particular argument is defended. These approaches are orthogonal to ours.

Our generalization of ADFs differs significantly from the mentioned papers in at least the following respects:

1. We are more general than existing work (with the exception of the work of [37]) in taking ADFs rather than AFs as starting point.
2. Rather than focusing on a single set of acceptance values, we provide a framework where the values can be freely selected based on the needs of a particular application.
3. Our semantics are a direct generalization of the operator-based ADF semantics and does not require the computation and aggregation of results for various subgraphs. Moreover, we obtain reasonable results also in cases where equational approaches do not have solutions.
4. Finally, the choice of an adequate information ordering allows us to do some fine tuning which is not possible in any approach we are aware of.

As to future work, we first want to explore restricted subclasses of weighted ADFs. In particular, we would like to exploit the known concept to express Dung AFs as ADFs (see [11],

Theorem 2) in order to investigate a new form of weighted AFs as a subclass of weighted ADFs. Our general definitions of the standard semantics for weighted ADFs will readily deliver natural definitions of semantics for such weighted AFs. Furthermore, the subclass of *bipolar* ADFs has been recognised as a useful class, as they are strictly more expressive than AFs while of equal computational complexity [41, 40, 36], so we intend to investigate *weighted bipolar* ADFs. A first step would be the generalization of supporting and attacking links to the multi-valued setting, for example via regarding of acceptance functions’ monotonicity and antimonicity in single (function) arguments [8].

We also would like to explore an idea that goes back to [10]. In our approach, as in standard AFT, interpretations $v : S \rightarrow V$ of atoms S with truth values V are approximated by functions $v' : S \rightarrow V_{\mathbf{u}}$ with $V_{\mathbf{u}} = V \cup \{\mathbf{u}\}$ for $\mathbf{u} \notin V$. Such three-valued/partial interpretations consequently represent the set of their completions. However, not all sets of total interpretations can be represented as completions of a partial interpretation. This is due to the fact that partial interpretations either assign a specific value, or leave the value completely undefined. This suggests the following: A *generalized partial interpretation (gpi)* of S in V is a total function $v : S \rightarrow 2^V \setminus \{\emptyset\}$, that is, a *gpi* assigns to each element of S a non-empty subset of the values in V . In this new setting, the total function $w : S \rightarrow V$ is a *completion of v* if and only if for all $s \in S$, we have $w(s) \in v(s)$. Based on the notion of a *gpi* we can generalize the characteristic ADF operator Γ_D to operate on *gpis* rather than partial interpretations. For each node s , the revised *gpi* $\Gamma_D(g)$ returns the set of values that are obtained by evaluating the acceptance condition of s under any completion of the input *gpi* g . A further investigation of this topic is on our agenda.

Acknowledgments

This work has been supported by the Deutsche Forschungsgemeinschaft (DFG) under grant BR 1817/7-2 and by the Austrian Science Fund (FWF): I2854 and P30168-N31.

References

- [1] Latifa Al-Abdulkarim, Katie Atkinson, and Trevor J. M. Bench-Capon. Abstract dialectical frameworks for legal reasoning. In *Proc. JURIX'14*, volume 271 of *FAIA*, pages 61–70. IOS Press, 2014.
- [2] Latifa Al-Abdulkarim, Katie Atkinson, and Trevor J. M. Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artif. Intell. Law*, 24(1):1–49, 2016.
- [3] Teresa Alsinet, Josep Argelich, Ramón Béjar, Cèsar Fernández, Carles Mateu, and Jordi Planes. Weighted argumentation for analysis of discussions in Twitter. *Int. J. Approx. Reasoning*, 85:21–35, 2017.
- [4] Leila Amgoud and Jonathan Ben-Naim. Evaluation of arguments in weighted bipolar graphs. In *Proc. ECSQARU'17*, volume 10369 of *LNCS*, pages 25–35. Springer, 2017.

- [5] Leila Amgoud, Jonathan Ben-Naim, Dragan Doder, and Srdjan Vesic. Acceptability semantics for weighted argumentation frameworks. In *Proc. IJCAI'17*, pages 56–62. IJCAI, 2017.
- [6] Leila Amgoud, Jonathan Ben-Naim, and Srdjan Vesic. Measuring the intensity of attacks in argumentation graphs with shapley value. In *Proc. IJCAI'17*, pages 63–69. IJCAI, 2017.
- [7] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Eng. Review*, 26(4):365–410, 2011.
- [8] Ringo Baumann and Hannes Strass. On the number of bipolar Boolean functions. *J. Log. Comput.*, 27(8):2431–2449, 2017.
- [9] Philippe Besnard and Anthony Hunter. A logic-based theory of deductive arguments. *Artif. Intell.*, 128(1-2):203–235, 2001.
- [10] Bart Bogaerts, Joost Vennekens, and Marc Denecker. On well-founded set-inductions and locally monotone operators. *ACM Trans. Comput. Log.*, 17(4):27:1–27:32, 2016.
- [11] Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes P. Wallner, and Stefan Woltran. Abstract Dialectical Frameworks Revisited. In *Proc. IJCAI'13*, pages 803–809. AAAI Press / IJCAI, 2013.
- [12] Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes P. Wallner, and Stefan Woltran. Abstract dialectical frameworks. An overview. *IfCoLog Journal of Logics and their Applications*, 4(8):2263–2318, 2017.
- [13] Gerhard Brewka and Stefan Woltran. Abstract Dialectical Frameworks. In *Proc. KR'10*, pages 102–111. AAAI Press, 2010.
- [14] Gerhard Brewka and Stefan Woltran. GRAPPA: A semantical framework for graph-based argument processing. In *Proc. ECAI'14*, volume 263 of *FAIA*, pages 153–158. IOS Press, 2014.
- [15] Elena Cabrio and Serena Villata. Abstract dialectical frameworks for text exploration. In *Proc. ICAART'16*, pages 85–95. SciTePress, 2016.
- [16] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Bipolar abstract argumentation systems. In Guillermo Ricardo Simari and Iyad Rahwan, editors, *Argumentation in Artificial Intelligence*, pages 65–84. Springer, 2009.
- [17] Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, second edition, 2002.
- [18] Marc Denecker, V. Wiktor Marek, and Mirosław Truszczyński. Uniform Semantic Treatment of Default and Autoepistemic Logics. *Artif. Intell.*, 143(1):79–122, 2003.
- [19] Marc Denecker, Victor W. Marek, and Mirosław Truszczyński. Approximations, Stable Operators, Well-Founded Fixpoints and Applications in Nonmonotonic Reasoning. In *Logic-Based Artificial Intelligence*, pages 127–144. Kluwer Academic Publishers, 2000.
- [20] Marc Denecker, Victor W. Marek, and Mirosław Truszczyński. Ultimate approximation and its application in nonmonotonic knowledge representation systems. *Inf. Comput.*, 192(1):84–121, 2004.

- [21] Pierpaolo Dondio. Multi-valued and probabilistic argumentation frameworks. In *Proc. COMMA'14*, volume 266 of *FAIA*, pages 253–260. IOS Press, 2014.
- [22] Pierpaolo Dondio. Propagating degrees of truth on an argumentation framework: an abstract account of fuzzy argumentation. In *Proc. SAC'17*, pages 995–1002. ACM, 2017.
- [23] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–357, 1995.
- [24] Phan Minh Dung and Phan Minh Thang. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proc. COMMA'10*, volume 216 of *FAIA*, pages 171–182. IOS Press, 2010.
- [25] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.*, 175(2):457–486, 2011.
- [26] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.*, 16(3):22:1–22:39, 2015.
- [27] Dov M. Gabbay. Equational approach to argumentation networks. *Argument & Computation*, 3(2-3):87–142, 2012.
- [28] Siegfried Gottwald. Many-valued logic. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2015 edition, 2015.
- [29] Davide Grossi and Sanjay Modgil. On the graded acceptability of arguments. In *Proc. IJCAI'15*, pages 868–874. AAAI Press, 2015.
- [30] Anthony Hunter. A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reasoning*, 54(1):47–81, 2013.
- [31] Anthony Hunter and Matthias Thimm. Probabilistic argument graphs for argumentation lotteries. In *Proc. COMMA'14*, volume 266 of *FAIA*, pages 313–324. IOS Press, 2014.
- [32] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with incomplete information. In *Proc. ECAI'14*, volume 263 of *FAIA*, pages 1033–1034. IOS Press, 2014.
- [33] Atefeh Keshavarzi. Investigating subclasses of abstract dialectical frameworks. Master's thesis, TU Wien, 2017.
- [34] João Leite and João Martins. Social abstract argumentation. In *Proc. IJCAI'11*, pages 2287–2292. IJCAI/AAAI, 2011.
- [35] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proc. TAFE'11, Revised Selected Papers*, volume 7132 of *LNCS*, pages 1–16. Springer, 2011.
- [36] Thomas Linsbichler, Jörg Pührer, and Hannes Strass. A Uniform Account of Realizability in Abstract Argumentation. In *Proc. ECAI'16*, volume 285 of *FAIA*, pages 252–260. IOS Press, September 2016.
- [37] Sylwia Polberg and Dragan Doder. Probabilistic abstract dialectical frameworks. In *Proc. JELIA'14*, volume 8761 of *LNCS*, pages 591–599. Springer, 2014.

- [38] Jörg Pührer. ArgueApply: A mobile app for argumentation. In *Proc. LPNMR'17*, volume 10377 of *LNCS*, pages 250–262. Springer, 2017.
- [39] Hannes Strass. Approximating Operators and Semantics for Abstract Dialectical Frameworks. *Artif. Intell.*, 205:39–70, 2013.
- [40] Hannes Strass. Expressiveness of Two-Valued Semantics for Abstract Dialectical Frameworks. *J. Artif. Intell. Res.*, 54:193–231, 2015.
- [41] Hannes Strass and Johannes P. Wallner. Analyzing the Computational Complexity of Abstract Dialectical Frameworks via Approximation Fixpoint Theory. *Artif. Intell.*, 226:34–74, 2015.