The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Weighted Abstract Dialectical Frameworks

Gerhard Brewka, Hannes Strass
Leipzig University, Computer Science Institute
Leipzig, Germany

Johannes P. Wallner, Stefan Woltran
TU Wien, Institute of Information Systems
Vienna, Austria

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.

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) (Dung 1995) 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. (Baroni, Caminada, and Giacomin 2011) 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 Cayrol and Lagasquie-Schiex, 2009). One of the most systematic and flexible outcomes of this research are abstract dialectical frameworks (ADFs) (Brewka and Woltran 2010; Brewka et al. 2013). 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 (Al-Abdulkarim, Atkinson, and Bench-Capon 2014; 2016) and text exploration (Cabrio and Villata 2016). A mobile argumentation app based on ADF techniques was developed by Pührer (2017).

The operator-based semantics of ADFs can be traced back to the work of Denecker, Marek, and Truszczyński (2000; 2003; 2004) on approximation fixpoint theory (AFT), an algebraic framework for studying semantics of knowledge representation formalisms. We refer to the work of Strass (2013) 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. (Alsinet et al. 2017) 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 $C_s : (par(s) \to \{\mathbf{t}, \mathbf{f}\}) \to \{\mathbf{t}, \mathbf{f}\}$. 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 $w : par(s) \to \{\mathbf{t}, \mathbf{f}\}$, 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 (Brewka et al. 2013) these semantics generalize the corresponding notions defined for AFs. For $\sigma \in \{adm, com, prf\}$, $\sigma(D)$ denotes the set of all admissible (resp. complete, preferred) interpretations with respect to D.

Example 1. Given ADF D over $\{a, b\}$ with $\varphi_a = a \vee \neg b$, $\varphi_b = \neg a$, and $v_1 = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}\}$, $v_2 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}\}$, $v_3 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}\}$, $v_4 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}\}$, we get $adm(D) = \{v_1, v_2, v_3, v_4\}$, $com(D) = \{v_1, v_3, v_4\}$ (note $\Gamma_D(v_2) = v_3$, thus $v_2 \notin com(D)$), and $prf(D) = \{v_3, v_4\}$. $\diamond$

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 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,
- 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 $C_s : (par(s) \to V) \to V$,
- $(V_{\mathbf{u}}, \leq_i)$ – where $V_{\mathbf{u}} = V \cup \{\mathbf{u}\}$ – forms a complete partial order with least element $\mathbf{u}$.

[1] Slightly 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.

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 $\bigsqcap_i X \in V_{\mathbf{u}}$, and (3) each ascending chain $x_1 \leq_i x_2 \leq_i \ldots$ over $V_{\mathbf{u}}$ has a least upper bound in $V_{\mathbf{u}}$.

ADFs are a special case of wADFs with $V = \{\mathbf{t}, \mathbf{f}\}$ and the information ordering as defined in the background section. We provide a formal result in Theorem 5 below.

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 $w : par(s) \to V$, $C_s(w)$ 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.

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 2. 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:

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

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 $v : S \to V_{\mathbf{u}}$ 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. A completion of v (as in the standard case) is any interpretation w that behaves like v whenever v is defined and assigns an arbitrary value from V otherwise. More formally: $w(a) = v(a)$ if $v(a) \neq \mathbf{u}$ and $w(a) \in V$ if $v(a) = \mathbf{u}$. Note that each completion w is a total interpretation $w : S \to V$ satisfying $v \leq_i w$.

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 $\{C_s(w) \mid w \in [v]_c\}$, where the set $[v]_c$ contains all completions of v. With these specifications the rest is entirely analogous to the definitions for standard ADFs.

Definition 3. Let $D = (S, L, C, V, \leq_i)$ be a wADF and $v : S \to V_{\mathbf{u}}$. Applying Γ_D to v yields a new interpretation (the consensus over $[v]_c$) defined as

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

where $\bigsqcap_i$ denotes the greatest lower bound in $(V_{\mathbf{u}}, \leq_i)$.

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

Definition 4. An interpretation v of a wADF $D = (S, L, C, V, \leq_i)$ is

- a *model* of D iff $v(s) \neq \mathbf{u}$ for all $s \in S$ and $\Gamma_D(v) = v$.
  Intuition: the value of a node s in v is exactly the one required by the acceptance condition of S.
- *grounded* for D iff $v = lfp(\Gamma_D)$, 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.

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

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. The simplest way to do this is to show monotonicity of the operator Γ_D.

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

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

The pair $(\{v : S \to V_{\mathbf{u}}\}, \leq_i)$ then forms a CPO in which the characteristic operator Γ_D of wADFs is monotone.

Proposition 1. *The operator Γ_D is $\leq_i$-monotone, that is: for all interpretations $v, w : S \to V_{\mathbf{u}}$ we have that $v \leq_i w$ implies $\Gamma_D(v) \leq_i \Gamma_D(w)$.*

Existence of the least fixpoint of Γ_D then follows via the fixpoint theorem for monotone operators in complete partial orders (see, e.g., Davey and Priestley, 2002, Theorem 8.22).

The following result is a generalization of Theorem 25 of Dung (1995) and Theorem 1 by Brewka et al. (2013).

Theorem 2. *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.*

[2] This differs from approaches like (Amgoud and Ben-Naim 2017) which consider weight assignments as part of the input and is more in line with research in multi-valued logics.

3. *The complete interpretations for D form a complete meet-semilattice with respect to $\leq_i$.*

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

Theorem 3. *Let D be a weighted* ADF. *It holds that*

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

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

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.

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 Keshavarzi (2017).

Theorem 4. *For any acyclic wADF D with finite S, $com(D) = prf(D) = \{v\}$ with v the grounded interpretation of D.*

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 $\mathbf{t}$ and $\mathbf{f}$ can appear in models, $\mathbf{f}$ (false) can be assumed to hold (by default), whereas $\mathbf{t}$ (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 5. Let $D = (S, L, C, V, \leq_i)$ be a wADF. Let $v : S \to V$ 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 $D^v_W = (S^v_W, L^v_W, C^v_W, V, \leq_i)$ where

- $S^v_W = \{s \in S \mid v(s) \notin W\}$,
- $L^v_W = L \cap (S^v_W \times S^v_W)$,
- $C^v_W = (C'_s)_{s \in S^v_W}$ where C'_s is obtained from C_s by fixing the value of each parent s^* of s in D such that $s^* \notin S^v_W$ 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^v_W by their v-values. Now stable models can be defined as usual:

Definition 6. Let $D = (S, L, C, V, \leq_i)$ be a wADF and let $v : S \to V$ 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 $s \in S^v_W$.

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 5. *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.*

4 Weighted ADFs Over the Unit Interval

In this section we focus on weighted ADFs over the unit interval, that is, wADFs over $V = [0, 1]$. 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 2. Consider wADF D over $[0, 1]$ depicted below.

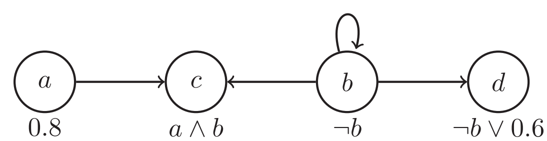

Intuitively, φ_a fixes the value of a to 0.8. φ_b expresses self-attack. φ_c means c is accepted to the extent a and b are, while d is attacked by b. In addition, it is known, for whatever reason, that the value of d must be at least 0.6.

We represent a partial interpretation v as the tuple $(v(a), v(b), v(c), v(d))$. 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 $W = \{x \in [0, 1] \mid x \leq 0.5\}$. To see this, consider the 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 W-stable for $W = \{x \in [0, 1] \mid x < 0.5\}$. 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 $v(c) = 0.6$ (if $v(b) = 0.5$).

The single preferred interpretation for D is v_2. ◇

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

Proposition 6. *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 $\Gamma_D(v)(s) = \mathbf{u}$.*

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 3. 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 $x \geq 0.5$. ◇

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

$$\textstyle\bigsqcap_i' X = \begin{cases} \mathbf{u} & \text{if } \mathbf{u} \in X \\ \inf X & \text{if } X \subseteq [0.5, 1] \\ \sup X & \text{if } X \subseteq [0, 0.5] \\ 0.5 & \text{otherwise} \end{cases}$$

where $\inf$ and $\sup$ are greatest lower bound and least upper bound in the complete lattice $([0, 1], \leq)$. 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 $lfp(\Gamma_D)$ exists and is uniquely determined.

Example 4. Consider again Example 2, 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 least fixpoint $v_1 = (0.8, 0.5, 0.5, 0.6)$. The more refined information ordering thus leads to more informative results, as expected. Due to space constraints determining the other semantics is left to the reader. ◇

In general, arbitrary CPO-preserving refinements of the information ordering have the following effects: (1) The model semantics is unaffected, since it is independent of the information ordering, but only depends on the truth values V; (2–4) grounded, admissible, preferred, and complete semantics may increase in information content, but never decrease.

Proposition 7. *Let $D = (S, L, C, V, \leq_i)$ be a wADF with information-ordering CPO $(V_\mathbf{u}, \leq_i)$. Furthermore let $\leq_i' \subseteq V_\mathbf{u} \times V_\mathbf{u}$ with $\leq_i \subseteq \leq_i'$ be such that $(V_\mathbf{u}, \leq_i')$ is a CPO, and define $D' = (S, L, C, V, \leq_i')$. Then for each $v : S \to V_\mathbf{u}$ and each $w : S \to V$, we have:*

1. *$\Gamma_D(w) = w$ if and only if $\Gamma_{D'}(w) = w$;*
2. *$lfp(\Gamma_D) \leq_i lfp(\Gamma_{D'})$;*
3. *$adm(\Gamma_D) \subseteq adm(\Gamma_{D'})$; and*
4. *if $v \in prf(\Gamma_D)$ ($v \in com(\Gamma_D)$) then there is a $v' \in prf(\Gamma_{D'})$ ($v' \in com(\Gamma_{D'})$) with $v \leq_i' v'$.*

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 Gottwald, 2015), e.g.

- $W_m = \{\frac{k}{m-1} \mid 0 \leq k \leq m-1\}$,[3] or
- Belnap's 4-valued system with $\{\emptyset, \{\bot\}, \{\top\}, \{\bot, \top\}\}$.

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

[3]As an example consider $W_5 = \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$.

W_3: Let us start with $W_3 = \{0, 0.5, 1\}$ 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.

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

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

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 $\{\bot\}$ and $\{\top\}$, and leaves the other two values unchanged. The truth ordering is the reflexive closure of:

$$\{\bot\} \leq_t \emptyset \leq_t \{\top\} \qquad \{\bot\} \leq_t \{\bot, \top\} \leq_t \{\top\}$$

The information ordering is the reflexive closure of:

$$\mathbf{u} \leq_i \emptyset \leq_i \{\top\} \leq_i \{\bot, \top\} \qquad \mathbf{u} \leq_i \emptyset \leq_i \{\bot\} \leq_i \{\bot, \top\}$$

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

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

Formula evaluation can be defined as follows:

$$[a, b] \wedge^* [c, d] = [min(a, c), min(b, d)]$$
$$[a, b] \vee^* [c, d] = [max(a, c), max(b, d)]$$
$$\neg^*[a, b] = [1 - b, 1 - a]$$

The information ordering is the $\subseteq$-least relation satisfying (1) $\mathbf{u} \leq_i v$ for all $v \in \mathit{INT}$, 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 (Strass and Wallner 2015) 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 8. *Verifying whether a given interpretation is admissible in wADFs with fixed and finite valuation structures is in* coNP.

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 9. *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.

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 10. *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.

Analogously to ADFs, the same complexity upper bound for existence of stable models can be derived for wADFs with fixed and finite valuation structures. The assumed values can be arbitrarily chosen among the fixed V.

Proposition 11. *Checking existence of stable models is in* Σ_2^p *for wADFs with fixed and finite valuation structures.*

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

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 (Brewka and Woltran 2010). 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 (Brewka and Woltran 2014). In (Dunne et al. 2011) 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 (Gabbay 2012). 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. (Dung and Thang 2010; Li, Oren, and Norman 2011; Hunter 2013; Hunter and Thimm 2014a; 2014b) (for a complexity analysis for probabilistic AFs see (Fazzinga, Flesca, and Parisi 2015)), and even of ADFs (Polberg and Doder 2014). 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 (Dondio 2014; 2017).

Social AFs (Leite and Martins 2011) 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 (Amgoud, Ben-Naim, and Vesic 2017; Amgoud et al. 2017). 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 (Amgoud et al. 2017) adapted to our setting are of interest and thus are on our agenda for future work.

Finally in (Besnard and Hunter 2001) 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 (2015) 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 Polberg and Doder, 2014) 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 (Brewka et al. 2013), 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 (Strass and Wallner 2015; Strass 2015; Linsbichler, Pührer, and Strass 2016), 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 antimonotonicity in single (function) arguments (Baumann and Strass 2017).

We also would like to explore an idea that goes back to Bogaerts, Vennekens, and Denecker (2016). In our approach, as in standard AFT, interpretations $v : S \to V$ of atoms S with truth values V are approximated by functions $v' : S \to V_{\mathbf{u}}$ 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 $v : S \to 2^V \setminus \{\emptyset\}$, 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 $w : S \to V$ 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.

Acknowledgements

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

Al-Abdulkarim, L.; Atkinson, K.; and Bench-Capon, T. J. M. 2014. Abstract dialectical frameworks for legal reasoning. In *Proc. JURIX'14*, volume 271 of *FAIA*, 61–70. IOS Press.

Al-Abdulkarim, L.; Atkinson, K.; and Bench-Capon, T. J. M. 2016. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artif. Intell. Law* 24(1):1–49.

Alsinet, T.; Argelich, J.; Béjar, R.; Fernández, C.; Mateu, C.; and Planes, J. 2017. Weighted argumentation for analysis of discussions in Twitter. *Int. J. Approx. Reasoning* 85:21–35.

Amgoud, L., and Ben-Naim, J. 2017. Evaluation of arguments in weighted bipolar graphs. In *Proc. ECSQARU'17*, volume 10369 of *LNCS*, 25–35. Springer.

Amgoud, L.; Ben-Naim, J.; and Vesic, S. 2017. Measuring the intensity of attacks in argumentation graphs with shapley value. In *Proc. IJCAI'17*, 63–69. IJCAI.

Amgoud, L.; Ben-Naim, J.; Doder, D.; and Vesic, S. 2017. Acceptability semantics for weighted argumentation frameworks. In *Proc. IJCAI'17*, 56–62. IJCAI.

Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *Knowledge Eng. Review* 26(4):365–410.

Baumann, R., and Strass, H. 2017. On the number of bipolar Boolean functions. *J. Log. Comp.* Advance Access Online 07 August 2017. http://dx.doi.org/10.1093/logcom/exx025.

Besnard, P., and Hunter, A. 2001. A logic-based theory of deductive arguments. *Artif. Intell.* 128(1-2):203–235.

Bogaerts, B.; Vennekens, J.; and Denecker, M. 2016. On well-founded set-inductions and locally monotone operators. *ACM Trans. Comput. Log.* 17(4):27:1–27:32.

Brewka, G., and Woltran, S. 2010. Abstract Dialectical Frameworks. In *Proc. KR'10*, 102–111. AAAI Press.

Brewka, G., and Woltran, S. 2014. GRAPPA: A semantical framework for graph-based argument processing. In *Proc. ECAI'14*, volume 263 of *FAIA*, 153–158. IOS Press.

Brewka, G.; Ellmauthaler, S.; Strass, H.; Wallner, J. P.; and Woltran, S. 2013. Abstract Dialectical Frameworks Revisited. In *Proc. IJCAI'13*, 803–809. AAAI Press / IJCAI.

Cabrio, E., and Villata, S. 2016. Abstract dialectical frameworks for text exploration. In *Proc. ICAART'16*, 85–95. SciTePress.

Cayrol, C., and Lagasquie-Schiex, M. 2009. Bipolar abstract argumentation systems. In *Argumentation in Artificial Intelligence*. Springer. 65–84.

Davey, B., and Priestley, H. 2002. *Introduction to Lattices and Order*. Cambridge University Press, second edition.

Denecker, M.; Marek, V. W.; and Truszczyński, M. 2000. Approximations, Stable Operators, Well-Founded Fixpoints and Applications in Nonmonotonic Reasoning. In *Logic-Based Artificial Intelligence*. Kluwer Academic Publishers. 127–144.

Denecker, M.; Marek, V. W.; and Truszczyński, M. 2003. Uniform Semantic Treatment of Default and Autoepistemic Logics. *Artif. Intell.* 143(1):79–122.

Denecker, M.; Marek, V. W.; and Truszczyński, M. 2004. Ultimate approximation and its application in nonmonotonic knowledge representation systems. *Inf. Comput.* 192(1):84–121.

Dondio, P. 2014. Multi-valued and probabilistic argumentation frameworks. In *Proc. COMMA'14*, volume 266 of *FAIA*, 253–260. IOS Press.

Dondio, P. 2017. Propagating degrees of truth on an argumentation framework: an abstract account of fuzzy argumentation. In *Proc. SAC'17*, 995–1002. ACM.

Dung, P. M., and Thang, P. M. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proc. COMMA'10*, volume 216 of *FAIA*, 171–182. IOS Press.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77(2):321–357.

Dunne, P. E.; Hunter, A.; McBurney, P.; Parsons, S.; and Wooldridge, M. 2011. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.* 175(2):457–486.

Fazzinga, B.; Flesca, S.; and Parisi, F. 2015. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.* 16(3):22:1–22:39.

Gabbay, D. M. 2012. Equational approach to argumentation networks. *Argument & Computation* 3(2-3):87–142.

Gottwald, S. 2015. Many-valued logic. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2015 edition.

Grossi, D., and Modgil, S. 2015. On the graded acceptability of arguments. In *Proc. IJCAI'15*, 868–874. AAAI Press.

Hunter, A., and Thimm, M. 2014a. Probabilistic argument graphs for argumentation lotteries. In *Proc. COMMA'14*, volume 266 of *FAIA*, 313–324. IOS Press.

Hunter, A., and Thimm, M. 2014b. Probabilistic argumentation with incomplete information. In *Proc. ECAI'14*, volume 263 of *FAIA*, 1033–1034. IOS Press.

Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reasoning* 54(1):47–81.

Keshavarzi, A. 2017. Investigating subclasses of abstract dialectical frameworks. Master's thesis, TU Wien.

Leite, J., and Martins, J. 2011. Social abstract argumentation. In *Proc. IJCAI'11*, 2287–2292. IJCAI/AAAI.

Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proc. TAFA'11, Revised Selected Papers*, volume 7132 of *LNCS*, 1–16. Springer.

Linsbichler, T.; Pührer, J.; and Strass, H. 2016. A Uniform Account of Realizability in Abstract Argumentation. In *Proc. ECAI'16*, volume 285 of *FAIA*, 252–260. IOS Press.

Polberg, S., and Doder, D. 2014. Probabilistic abstract dialectical frameworks. In *Proc. JELIA'14*, volume 8761 of *LNCS*, 591–599. Springer.

Pührer, J. 2017. ArgueApply: A mobile app for argumentation. In *Proc. LPNMR'17*, volume 10377 of *LNCS*, 250–262. Springer.

Strass, H., and Wallner, J. P. 2015. Analyzing the Computational Complexity of Abstract Dialectical Frameworks via Approximation Fixpoint Theory. *Artif. Intell.* 226:34–74.

Strass, H. 2013. Approximating Operators and Semantics for Abstract Dialectical Frameworks. *Artif. Intell.* 205:39–70.

Strass, H. 2015. Expressiveness of Two-Valued Semantics for Abstract Dialectical Frameworks. *J. Artif. Intell. Res.* 54:193–231.