

Conditional Abstract Dialectical Frameworks

Citation for published version (APA):

Heyninck, J., Thimm, M., Kern-Isberner, G., Rienstra, T., & Skiba, K. (2022). Conditional Abstract Dialectical Frameworks. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI'22): AAAI Technical Track on Knowledge Representation and Reasoning* (5 ed., Vol. 36, pp. 5692-5699). The Association for the Advancement of Artificial Intelligence .
<https://doi.org/10.1609/aaai.v36i5.20511>

Document status and date:

Published: 01/01/2022

DOI:

[10.1609/aaai.v36i5.20511](https://doi.org/10.1609/aaai.v36i5.20511)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Conditional Abstract Dialectical Frameworks

Jesse Heyninck^{1,2}, Matthias Thimm³, Gabriele Kern-Isberner¹, Tjitze Rienstra⁴, Kenneth Skiba³

¹Technische Universität Dortmund, Dortmund, Germany

²University of Cape Town and CAIR, South-Africa

³FernUniversität Hagen, Hagen, Germany

⁴Maastricht University, Maastricht, The Netherlands

Abstract

Abstract dialectical frameworks (in short, ADFs) are a unifying model of formal argumentation, where argumentative relations between arguments are represented by assigning acceptance conditions to atomic arguments. This idea is generalized by letting acceptance conditions being assigned to complex formulas, resulting in conditional abstract dialectical frameworks (in short, cADFs). We define the semantics of cADFs in terms of a non-truth-functional four-valued logic, and study the semantics in-depth, by showing existence results and proving that all semantics are generalizations of the corresponding semantics for ADFs.

1 Introduction

Formal argumentation is one of the major approaches to knowledge representation. In the seminal paper (Dung 1995), *abstract argumentation frameworks* were conceived of as directed graphs where nodes represent arguments and edges between these nodes represent attacks. So-called *argumentation semantics* determine which sets of arguments can be reasonably upheld together given such an argumentation graph. Various authors have remarked that other relations between arguments are worth consideration. For example, in (Cayrol and Lagasque-Schiex 2005), *bipolar argumentation frameworks* are developed, where arguments can support as well as attack each other.

The last decades saw a proliferation of such extensions of the original formalism of (Dung 1995), and it has often proven hard to compare the resulting different dialects of the argumentation formalisms. To cope with the resulting multiplicity, (Brewka et al. 2013) introduced *abstract dialectical frameworks* (in short, ADFs) that aims to unify these different dialects (Polberg 2016). Just like in (Dung 1995), ADFs are directed graphs. In difference to abstract argumentation frameworks, however, in ADFs, edges between nodes do not necessarily represent attacks but can encode *any* relationship between arguments. Such a generality is achieved by associating an *acceptance condition* with each argument, which is a Boolean formula in terms of the parents of the argument that expresses the conditions under which an argument can be accepted. This results in an ADF being defined as a triple

(At, L, C) where At represents a set of atoms or arguments, $L \subseteq At \times At$ represents a set of argumentative relations between the atoms and C is a set of acceptance conditions C_s for every $s \in At$. As such, ADFs are able to capture all of the major semantics of abstract argumentation and offer a general framework for argumentation-based inference. Furthermore, ADFs were shown to capture logic programming (Brewka et al. 2013). In (Heyninck et al. 2019), first attempts were made to translate non-monotonic conditional logics in ADFs.

However, there are limits to the representative capabilities of ADFs, both on a conceptual as well as a more technical level. On the conceptual level, acceptance conditions are assigned to atoms, which means that, e. g., an attack on a set of arguments cannot be captured by ADFs. For example, to state that the set $\{p, q\}$ is attacked by r we would have to be able to set the acceptance condition of $p \wedge q$ to $\neg r$, which is not possible in ADFs. Likewise, it is not immediately obvious how to represent more complicated logic programming languages in ADFs, such as disjunctive logic programming. Such limitations are, not unsurprisingly, also reflected on a more technical level. For example, a (polynomial) translation of disjunctive logic programming into ADFs is impossible in view of considerations on complexity. Finally, in (Heyninck et al. 2019) it was shown that only a fragment of the full language of conditional logics can be translated in ADFs in view of their limited syntax.

In this paper, we generalize ADFs as to allow for the assignment of acceptance conditions to complex formulas. This results in *conditional abstract dialectical frameworks* (in short, cADFs) which are sets of acceptance pairs of the form $\phi \triangleleft \psi$ with arbitrary formulas ϕ and ψ , interpreted as a defeasible version of “ ϕ is the case if and only if ψ is the case”. The semantics of cADFs are formulated as a generalization of the semantics of ADFs, with the Γ -function, on its turn based on a non-truth-functional four-valued logic, as a central component. Some of the main results include existence results for all the major semantics, as well as the definition of the so-called *grounded state*, a single-state semantics which can be iteratively constructed and represents the minimal information entailed by a given cADF.

Outline of this Paper: We first state all the necessary preliminaries in Section 2 on propositional logic (Section 2), and abstract dialectical argumentation (Section 2). The syn-

tax of *conditional abstract dialectical frameworks* cADFs is introduced in Section 3. In Section 4, a four-valued logic, which will form the basis of the semantics of cADFs, is defined and studied. In Section 5, we then define and study the admissible, complete, preferred and grounded semantics for cADFs. A unique, iteratively constructible analogue to the grounded extension, called the *grounded state*, is introduced in Section 6. Related work is discussed in Section 7 and a conclusion is drawn in Section 8.

2 Preliminaries

In the following, we briefly recall some general preliminaries on propositional logic and ADFs (Brewka et al. 2013).

Propositional Logic

For a set At of atoms let $\mathcal{L}(\text{At})$ be the corresponding propositional language constructed using the usual connectives \wedge (*and*), \vee (*or*), \neg (*negation*) and \rightarrow (*material implication*). A (classical) *interpretation* (also called *possible world*) ω for a propositional language $\mathcal{L}(\text{At})$ is a function $\omega : \text{At} \rightarrow \{\text{T}, \text{F}\}$. Let $\mathcal{V}^2(\text{At})$ denote the set of all interpretations for At . We simply write \mathcal{V}^2 if the set of atoms is implicitly given. An interpretation ω *satisfies* (or is a *model* of) an atom $a \in \text{At}$, denoted by $\omega \models a$, if and only if $\omega(a) = \text{T}$. The satisfaction relation \models is extended to formulas as usual. For $\Phi \subseteq \mathcal{L}(\text{At})$ we also define $\omega \models \Phi$ if and only if $\omega \models \phi$ for every $\phi \in \Phi$. Define the set of models $\text{Mod}^2(X) = \{\omega \in \mathcal{V}^2(\text{At}) \mid \omega \models X\}$ for every formula or set of formulas X . A formula or set of formulas X_1 *entails* another formula or set of formulas X_2 , denoted by $X_1 \vdash X_2$, if $\text{Mod}^2(X_1) \subseteq \text{Mod}^2(X_2)$. A formula ϕ is a tautology if $\text{Mod}^2(\phi) = \mathcal{V}^2(\text{At})$ and a falsity if $\text{Mod}^2(\phi) = \emptyset$.

Abstract Dialectical Frameworks

We briefly recall some technical details on ADFs following loosely the notation from (Brewka et al. 2013). An ADF D is a tuple $D = (\text{At}, L, C)$ where At is a finite set of atoms, $L \subseteq \text{At} \times \text{At}$ is a set of links, and $C = \{C_s\}_{s \in \text{At}}$ is a set of total functions $C_s : 2^{\text{par}_D(\text{At})} \rightarrow \{\text{T}, \perp\}$ for each $s \in \text{At}$ with $\text{par}_D(s) = \{s' \in \text{At} \mid (s', s) \in L\}$ (also called acceptance functions). An acceptance function C_s defines the cases when the statement s can be accepted (truth value T), depending on the acceptance status of its parents in D . By abuse of notation, we will often identify an acceptance function C_s by its equivalent *acceptance condition* which models the acceptable cases as a propositional formula.

Example 1. We consider the following ADF $D_1 = (\{a, b, c\}, L, C)$ with $L = \{(a, b), (b, a), (a, c), (b, c)\}$ and $C_a = \neg b$, $C_b = \neg a$, and $C_c = \neg a \vee \neg b$.

Informally, the acceptance conditions can be read as “ a is accepted if b is not accepted”, “ b is accepted if a is not accepted” and “ c is accepted if a or b is not accepted”.

An ADF $D = (\text{At}, L, C)$ is interpreted through 3-valued interpretations $\nu : \text{At} \rightarrow \{\text{T}, \text{F}, \text{U}\}$. We denote the set of all 3-valued interpretations over At by $\mathcal{V}^3(\text{At})$. We define the information order $<_i$ over $\{\text{T}, \text{F}, \text{U}\}$ by making U the minimal element: $\text{U} <_i \text{T}$ and $\text{U} <_i \text{F}$, and $\dagger \leq_i \ddagger$ iff

$\dagger <_i \ddagger$ or $\dagger = \ddagger$ for any $\dagger, \ddagger \in \{\text{T}, \text{F}, \text{U}\}$. This order is lifted point-wise as follows (given $\nu, \nu' \in \mathcal{V}^3(\text{At})$): $\nu \leq_i \nu'$ iff $\nu(s) \leq_i \nu'(s)$ for every $s \in \text{At}$. The set of two-valued interpretations extending a 3-valued interpretation ν is defined as $[\nu]^2 = \{\omega \in \mathcal{V}^2(\text{At}) \mid \nu \leq_i \omega\}$. Given a set of 3-valued interpretations $V \subseteq \mathcal{V}^3(\text{At})$, $\sqcap_i V$ is the 3-valued interpretation defined via $\sqcap_i V(s) = \dagger$ if for every $\nu \in V$, $\nu(s) = \dagger$, for any $\dagger \in \{\text{T}, \text{F}, \text{U}\}$, and $\sqcap_i V(s) = \text{U}$ otherwise. Truth values based on a three-valued interpretations can now be assigned to complex formulas ϕ by taking $\sqcap_i [\nu]^2(\phi)$. All major semantics of ADFs single out three-valued interpretations in which the truth value of every atom $s \in \text{At}$ is, in some sense, in alignment or agreement with the truth value of the corresponding condition C_s . The Γ -function enforces this intuition by mapping an interpretation ν to a new interpretation $\Gamma_D(\nu)$, which assigns to every atom s exactly the truth value assigned by ν to C_s , i.e.:

$$\Gamma_D(\nu) : \text{At} \mapsto \{\text{T}, \text{F}, \text{U}\} \text{ where } s \mapsto \sqcap_i \{\omega(C_s) \mid \omega \in [\nu]^2\}.$$

Definition 1. Let $D = (\text{At}, L, C)$ be an ADF with $\nu \in \mathcal{V}(\text{At})$ a 3-valued interpretation. Then: ν is *admissible* for D iff $\nu \leq_i \Gamma_D(\nu)$; ν is *complete* for D iff $\nu = \Gamma_D(\nu)$; ν is *preferred* for D iff ν is \leq_i -maximal among all admissible interpretations; and ν is *grounded* for D iff ν is \leq_i -minimal among all complete interpretations. We denote by $\text{adm}(D)$, $\text{cmp}(D)$, $\text{prf}(D)$, and $\text{grnd}(D)$ the sets of complete, preferred, and grounded interpretations of D , respectively.

Notice that ν is admissible iff $\nu(s) \leq_i \sqcap_i [\nu]^2(C_s)$ for every $s \in S$ and likewise, ν is complete iff $\nu(s) = \sqcap_i [\nu]^2(C_s)$ for every $s \in S$. It can thus be observed that the logic defined by $\sqcap_i [\nu]^2$ is, essentially, the logic underlying ADFs, in the sense that the evaluation of acceptance conditions under $\sqcap_i [\nu]^2$ is the fundamental operation underlying every semantical notion of ADFs. It should be furthermore noted that $\sqcap_i [\nu]^2$ does not give rise to a *truth-functional logic*. Recall that a truth-functional logic is a logic in which the truth value assigned to a complex formula is a function of the truth values of its component formulas. E.g. for a truth-functional logic, the truth value of $a \vee \neg b$ is determined completely by the truth value of a and $\neg b$. For example, given $\nu(a) = \text{U}$ and $\nu(b) = \text{U}$, $\sqcap_i [\nu]^2(\neg a) = \text{U}$ and $\sqcap_i [\nu]^2(a \vee \neg a) = \text{T}$ whereas $\sqcap_i [\nu]^2(a \vee \neg b) = \text{U}$. Thus, the logic defined by $\sqcap_i [\nu]^2$ is not truth-functional.

Example 2 (Ex. 1 ctd.). The ADF in Ex. 1 has three complete models ν_1, ν_2, ν_3 with: $\nu_1(a) = \text{T}, \nu_1(b) = \text{F}$ and $\nu_1(c) = \text{T}$; $\nu_2(a) = \text{F}, \nu_2(b) = \text{T}$ and $\nu_2(c) = \text{T}$; and $\nu_3(a) = \text{U}, \nu_3(b) = \text{U}$, and $\nu_3(c) = \text{U}$. ν_3 is the grounded interpretation whereas ν_1 and ν_2 are both preferred.

3 Syntax of cADFs

The syntactical representation $D = (S, L, C)$ of an ADF contains some superfluous information. In particular, as there is a link between a statement s and s' iff s is mentioned in the acceptance condition of s' , the set of links does not contain any information not already derivable from the set of acceptance conditions C . As such, given a set of atoms S , we can simply write an ADF as a set of statements $s \triangleleft C_s$

if C_s is the acceptance condition of s . So the ADF D_1 from Example 1 can be simply written as:

$$D_1 = \{a \triangleleft \neg b, b \triangleleft \neg a, c \triangleleft \neg a \vee \neg b\}$$

An ADF is determined by a set of propositional formulæ that, when evaluated to true, make a certain statement, which is a simple atom, true as well, and when evaluated to false, make the simple atom false as well. In other words, \triangleleft can be read as a *approximate if and only if*: $s \triangleleft C_s$ means that the truth-values s and C_s should be aligned. \triangleleft can truly be read as a *approximate iff*, since it might not always be possible to align the truth values of s and C_s in such a way that they take on exactly the same (determinate) truth value. To see this, consider, e. g., $a \triangleleft \neg a$. We generalise this framework by allowing these statements to be arbitrary propositional formulæ:

Definition 2. Given a set of atoms At , a *conditional abstract dialectical framework cADF* Π w.r.t. At is a finite set of *acceptance pairs* over At , where an *acceptance pair* is of the form:

$$\phi \triangleleft \psi$$

with ϕ and ψ being propositional formulæ over At .

In order to stick to ADF terminology we call ϕ the *statement* and ψ the *condition* of the acceptance pair $\phi \triangleleft \psi$. We omit the reference to the signature At when it is clear from context.

Example 3. Consider a cADF $\Pi_1 = \{c_1, c_2, c_3\}$ with

$$\begin{aligned} c_1 : & \quad p \vee s \vee q \triangleleft \top \\ c_2 : & \quad p \wedge s \triangleleft \neg q \\ c_3 : & \quad (p \wedge q) \vee (p \wedge s) \triangleleft t \end{aligned}$$

This cADF can be used to model an argument of a group of friends about making plans on a Sunday. They are discussing whether to go to a party (p), to the swimming pool (s) or go to a pub quiz (q). They want to do at least one of these three things (c_1). However, if they go to the quiz, they won't be able to still go to the pool *and* go to the party (represented by the attack of q on $p \wedge s$ in c_2). If everyone arrives on time (t), they would like to go to both the quiz and the party, or to both the pool and the party (c_3). We notice that without adding further atoms, an attack from q on the set $\{p, s\}$, as formalized by c_2 , cannot be represented in ADFs.

We observe that this simple generalization w.r.t. ADFs results in the following additional points of expressiveness in comparison to ADFs:

- cADFs allow for complex formulas as statements, as demonstrated by $(p \wedge q) \vee (p \wedge s) \triangleleft t$ in Example 3
- cADFs allow for “incomplete” specifications, i.e. they do not force the user to formulate an acceptance condition for every atom, as demonstrated by the cADF $\{a \triangleleft b\}$, where b has no acceptance condition.
- cADFs allow for “overspecifications” or conflicting specifications, as demonstrated by the cADF $\{a \triangleleft b, \neg a \triangleleft b\}$ where both a and $\neg a$ have the acceptance condition b .


- cADFs allow for indeterminism, as demonstrated by the cADF $\{a \vee b \triangleleft \top\}$, where $a \vee b$ is required to be true, but no further information on which of the disjuncts is required to be true is given.

To cope with this higher expressiveness semantically, it will prove useful to move from three-valued interpretations to four-valued interpretations. To assign truth values to complex formulas on the basis of four-valued interpretations, we generalize the logic defined by $\sqcap_i[\nu]^2$ to a four-valued setting in Section 4. We then generalize the semantics of ADFs to cADFs on the basis of this four-valued logic in Section 5.

4 A Four-Valued Logic Based on Completions

We first define a four-valued logic 4CM which generalizes the idea of completions known from the logic underlying ADFs defined by $[\nu]^2$, which preserves classical tautologies and falsities. We first recall four-valued interpretations. A *four-valued interpretation* $v : \text{At} \rightarrow \{\top, \text{F}, \text{I}, \text{U}\}$ assigns to every atom a truth value \top (true), F (false), U (undecided) or I (inconsistent). We will also write an interpretation $v \in \mathcal{V}^4(\{a_1, \dots, a_n\})$ as $v(a_1) \dots v(a_n)$, e. g., v over $\{p, q\}$ with $v(p) = \top$ and $v(q) = \text{U}$ will be written as TU . We denote the set of four-valued interpretations over At by $\mathcal{V}^4(\text{At})$. Notice that $\mathcal{V}^2(\text{At}) \subseteq \mathcal{V}^3(\text{At}) \subseteq \mathcal{V}^4(\text{At})$. If it is clear that an interpretation is two- respectively three-valued, we will denote it by (a possibly indexed) ω respectively ν .

Two useful orders over these truth values are the *information order* \leq_i and the *truth order* \leq_t , which form the following bilattice-structure (Belnap 1977):



Notice that $\mathcal{V}^4(\text{At})$ also forms a bounded lattice under \leq_i with v_{U} and v_{I} as least and greatest element respectively (where v_{U} is defined as the interpretation that sets $v_{\text{U}}(a) = \text{U}$ for every $a \in \text{At}$ and v_{I} is defined as $v_{\text{I}}(a) = \text{I}$ for every $a \in \text{At}$).

We shall interpret the four truth values, at least for atoms, in the same way as (Belnap 2019): U (*undecided*) means that we have no explicit information for either the truth nor the falsity of an atom. T (*true*) respectively F (*false*) means that we have explicit information only for the truth respectively the falsity of the atom in question. Finally, I (*inconsistent*) means that we have explicit information for both the truth and the falsity of the atom in question. When it comes to complex formulas, we take a somewhat hybrid position between truth values expressing merely explicit information and truth values standing for objective truth. In particular, the logic we will define here will allow for *logically contingent* formulas, i. e., formulas which are neither classical tautologies nor classical falsities, to be assigned any of the four truth values, whereas classical tautologies and classical falsities will always be assigned \top respectively F by any

interpretation. Intuitively, this means that even though the truth value of $s \in \text{At}$ might be undetermined (U) or inconsistent (I), the logic will still evaluate $s \vee \neg s$ as true. This is in complete agreement with ADFs, where tautologies and logical falsities are always evaluated in agreement with classical logic by $\sqcap_i [v]^2$.

Semantically, we proceed as follows: we construct a set of sets of (two-valued) worlds on the basis of a four-valued interpretation v that represents the beliefs expressed by v . Just like in the logic underlying ADFs $\sqcap_i [v]^2$, a set of (two-valued) worlds will be used to represent a three-valued interpretation ν . The worlds in $[\nu]^2$ represent equally plausible candidates of the actual world in view of the beliefs expressed by the three-valued interpretation ν . Likewise, a set of three-valued interpretations $[\nu]^3$ will be used to represent the information expressed by a four-valued interpretation v . $[\nu]^3$ consists of the three-valued interpretations that *jointly* represent the information expressed by v . Notice the difference with $[\nu]^2$: $[\nu]^2$ consists of equally plausible candidates of the actual world in view of the information expressed by v , whereas $[\nu]^3$ contains interpretations that *taken together* represent the information expressed by v . We now develop this idea in more formal details.

Given a four-valued interpretation, we define the set of two-valued completions of v , $[v]^2$, in two steps. First, we construct $[\nu]^3$, which converts $v \in \mathcal{V}^4(\text{At})$ to a set of three-valued interpretations $[\nu]^3 \subseteq \mathcal{V}^3(\text{At})$. Then, we obtain $[v]^2 \subseteq \wp(\mathcal{V}^2(\text{At}))$ by converting every three-valued interpretation $\nu \in [\nu]^3$ to a set of two-valued interpretations $[\nu]^2$.

Definition 3. Given a four-valued interpretation $v \in \mathcal{V}^4(\text{At})$, $[\nu]^3 = \{\nu \in \mathcal{V}^3(\text{At}) \mid \text{for every } s \in \text{At} : \text{if } v(s) = \text{I then } \nu(s) \in \{\text{T}, \text{F}\}, \nu(s) = v(s) \text{ otherwise}\}$

In other words, $[\nu]^3$ is obtained by replacing every assignment of an atom s to I to an assignment of s to T or to F.

Notice that $[\nu]^3$ consists of the \leq_i -maximal three-valued interpretations that v extends:

Fact 1.¹ For any $v \in \mathcal{V}^4(\text{At})$, $[\nu]^3 = \max_{\leq_i}(\{\nu \in \mathcal{V}^3(\text{At}) \mid \nu \leq_i v\})$.

Example 4. Consider $v = \text{TUI}$ over $\Sigma = abc$. Then $[\nu]^3 = \{\text{TUT}, \text{TUF}\}$.

We are now ready to define the *four-valued completions* $[v]^4$ of v :

Definition 4. Given some $v \in \mathcal{V}^4(\text{At})$, the *four-valued completions* of v are defined as: $[v]^4 = \{[v']^2 \mid v' \in [\nu]^3\}$.

Thus, $[v]^4$ is obtained by first constructing $[\nu]^3$, and then taking for every $\nu \in [\nu]^3$ the set of two-valued completions of ν . The intuition behind this is as follows: $v(s) = \text{I}$ means that we have information for both s being true and s being false. Thus, the interpretations where we set $\nu_1(s) = \text{T}$ and $\nu_2(s) = \text{F}$ are both (partial yet consistent) representations of the state of the world represented by v . Hence $[\nu]^3$ can be viewed as the set of three-valued interpretations that together form the representation of the state of the world represented by v . We then construct for every such representation a set of two-valued interpretations, which represent equally

plausible candidates of the state of the world represented by $\nu \in [\nu]^3$. Altogether, $[v]^4$ contains a set of set of possible worlds, which together represent our knowledge about the actual state of the world.

It is useful to notice that for a three-valued interpretation $v \in \mathcal{V}^3(\text{At})$, $[v]^4 = \{[v]^2\}$.

Example 5. Consider $v = \text{TUI}$ over $\Sigma = \{abc\}$. Since $[\nu]^3 = \{\text{TUT}, \text{TUF}\}$, $[\text{TUT}]^2 = \{\text{TTT}, \text{TFT}\}$ and $[\text{TUF}]^2 = \{\text{TTF}, \text{TFF}\}$, we see that $[v]^4 = \{\{\text{TTT}, \text{TFT}\}, \{\text{TTF}, \text{TFF}\}\}$.

Notice that, in order to retain the four-valued structure of an interpretation v in its four-valued completion $[v]^4$, the two-step nature of the construction of $[v]^4$ and the resulting nested structure of $[v]^4$ is essential. Indeed, if $[v]^4$ would merely consist of possible worlds, we would somehow have to choose between letting the members $\omega \in [v]^4$ stand as equally plausible candidates of the actual world or partial descriptions of the information given by v , i. e., we would have to choose between U and I. Conceiving of $[v]^4$ as a set of sets of worlds avoids this choice: sets of worlds $\mathcal{V}' \in [v]^4$ represent partial descriptions of the information given by v , and members of these sets of worlds $\omega \in \mathcal{V}'$ represent equally plausible candidates of the information in \mathcal{V}' .

We can now define the assignment of truth values of complex formulas given an interpretation v based on our set of four-valued completions $[v]^4$:

Definition 5. Given a formula ϕ and an interpretation v , then:

$$v(\phi) = \begin{cases} \text{T} & \text{if for every } \Omega' \in [v]^4, \sqcap_i \Omega'(\phi) = \text{T} \\ \text{F} & \text{if for every } \Omega' \in [v]^4, \sqcap_i \Omega'(\phi) = \text{F} \\ \text{I} & \text{if for some } \Omega_1 \in [v]^4, \sqcap_i \Omega_1(\phi) = \text{T} \\ & \text{and for some } \Omega_2 \in [v]^4, \sqcap_i \Omega_2(\phi) = \text{F} \\ \text{U} & \text{otherwise} \end{cases}$$

Thus, a complex formula ϕ is assigned T (respectively F) relative to an interpretation v if every four-valued completion $\Omega' \in [v]^4$ of v , assigns T (respectively F) to ϕ . If there is disagreement among the four-valued completions of v on which determinate truth value ϕ should be assigned, $v(\phi) = \text{I}$. Finally, if some of the four-valued completions of v do not assign any determinate truth value to ϕ , $v(\phi) = \text{U}$.

This way of deriving a truth value for complex formulas on the basis of a four-valued interpretation is, to the best of our knowledge, completely new. It is perfectly in line with $\sqcap_i [v]^2$, the logic underlying ADFs, in the sense that for any three-valued interpretation $\nu \in \mathcal{V}^3(\text{At})$ and any formula $\phi \in \mathcal{L}$, $\nu(\phi) = \sqcap_i [\nu]^2(\phi)$.

Fact 2. For any $\nu \in \mathcal{V}^3(\text{At})$ and any $\phi \in \mathcal{L}(\text{At})$, $\nu(\phi) = \sqcap_i [\nu]^2(\phi)$.

Example 6. Consider $v = \text{TUI}$ over $\Sigma = abc$. Observe that $[v]^4 = \{\{\text{TTT}, \text{TFT}\}, \{\text{TTF}, \text{TFF}\}\}$. Thus, we have the following assignments to complex formulas:

- $v(a \wedge c) = \text{I}$, since $\sqcap_i \{\text{TTT}, \text{TFT}\}(a \wedge c) = \text{T}$ and $\sqcap_i \{\text{TTF}, \text{TFF}\}(a \wedge c) = \text{F}$;
- $v(b \wedge c) = \text{U}$, since $\sqcap_i \{\text{TTT}, \text{TFT}\}(b \wedge c) = \text{U}$ and $\sqcap_i \{\text{TTF}, \text{TFF}\}(b \wedge c) = \text{F}$;

¹Due to spatial limitations, some proofs are left out.

- $v(a \wedge \neg a) = F$, since $\sqcap_i\{TTT, TFT\}(a \wedge \neg a) = F$ and $\sqcap_i\{TTF, TFF\}(a \wedge \neg a) = F$;

Remark 1. Observe that the logic 4CM, like the logic defined by $\sqcap_i[v]^2$, is *not* truth-functional. To see this consider the interpretation v with $v(a) = U$ and $v(b) = U$. Then $v(a \vee \neg a) = T$ yet $v(b \vee \neg a) = U$. Thus, we see that 4CM is not truth-functional, as $v(a) = v(b) = U$ yet $v(a \vee \neg a) \neq v(b \vee \neg a)$.

5 Semantics of cADFs

In this section, we define, motivate and study the semantics of cADFs. We first define the central Γ_Π -function and use it to define the main semantics for cADFs. Then we motivate the design choices made in generalizing the Γ -function from ADFs to cADFs. Finally, we show semantic properties of the semantics of cADFs.

The Γ_Π -function and resulting cADF-semantics

A cADF Π over At is interpreted through 4-valued interpretations. Just like for ADFs, it is of crucial importance to construct a Γ -function that allows to characterize all semantics in terms of (post-)fixpoints of this function.

The Γ -function, conceptually, performs the following operation for ADFs: given an interpretation ν and an ADF D , $\Gamma_D(\nu)$ assigns to every atom s the truth value determined by ν and C_s . In other words, $\Gamma_D(\nu)(s)$ is the value s should take in view of the information expressed by $s \triangleleft C_s$ and ν . If (for every $s \in S$), this value is compatible (in terms of \leq_i) with the actual value $\nu(s)$, then ν will be admissible or even complete. We generalize this idea to the case of cADFs, and take, intuitively, $\Gamma_\Pi(v)$ as the *set of interpretations* that evaluate ϕ in accordance with the information given by $\phi \triangleleft \psi \in \Pi$ and v . More formally, we define the Γ -function $\Gamma_\Pi : \mathcal{V}^4(\text{At}) \rightarrow \wp(\mathcal{V}^4(\text{At}))$ for a cADF Π and an interpretation $v \in \mathcal{V}^4(\text{At})$ as follows:

$$\Gamma_\Pi(v) = \min_{\leq_i} \{v' \in \mathcal{V}^4 \mid \forall \phi \triangleleft \psi \in \Pi : v'(\phi) \geq_i v(\psi)\}$$

Example 7. Let $\Pi = \{p \vee s \triangleleft T; \neg s \triangleleft p\}$ formulated over the signature $\Sigma = \{p, s\}$. We have the following interpretations and corresponding outcomes of the Γ_Π -function:

| v | $\Gamma_\Pi(v)$ | v | $\Gamma_\Pi(v)$ |
|-----|-----------------|-----|-----------------|
| UU | {UT, TU} | FU | {UT} |
| UT | {UT, TU} | FT | {UT} |
| UF | {UT, TU} | FF | {UT} |
| UI | {UT, TU} | FI | {UT} |
| TU | {TF, FI} | IU | {TI, FI} |
| TT | {TF, FI} | IT | {TI, FI} |
| TF | {TF, FI} | IF | {TI, FI} |
| TI | {TF, FI} | II | {TI, FI} |

We explain $\Gamma_\Pi(UU)$ as follows: in view of $p \vee s \triangleleft T$ and $UU(T) = T$, every interpretation $v' \in \Gamma_\Pi(UU)$ has to assign a truth value at least as informative as T to $p \vee s$, i.e. $v'(p \vee s) \geq_i T$. Likewise, since $UU(p) = U$ and $\neg s \triangleleft p \in \Pi$, $v' \in \Gamma_\Pi(UU)$ has to set $v'(\neg s) \geq_i U$, which is trivially the case. The two \leq_i -minimal interpretations that satisfy this constraint are: UT and TU.

As a second example, consider FF. Like with UU, every interpretation $v' \in \Gamma_\Pi(FF)$ has to assign $v'(p \vee s) \geq_i T$. However, since $FF(p) = F$ and $\neg s \triangleleft p \in \Pi$, any $v' \in \Gamma_\Pi(FF)$ has to set $v'(\neg s) \geq_i F$. UT is the unique \leq_i -minimal interpretation satisfying these constraints.

We first notice that Γ_Π is indeed a generalization of the Γ_D -function for ADFs. To show this in a more formally precise manner, we first define the cADF Π_D associated with an ADF D .

Definition 6. Given an ADF $D = (S, L, C)$, we define the cADF Π_D associated with D as $\Pi_D = \{s \triangleleft C_s \mid s \in S\}$.

We can now show that for any three-valued interpretation ν , $\Gamma_{\Pi_D}(\nu)$ coincides with $\Gamma_D(\nu)$, i.e. the Γ -function for ADFs coincides with the Γ -function for the associated cADFs for three-valued interpretations.

Proposition 1. For any ADF $D = (S, L, C)$ and any $\nu \in \mathcal{V}^3(S)$, $\Gamma_{\Pi_D}(\nu) = \{\Gamma_D(\nu)\}$.

The above result shows that the Γ_Π -function is a direct generalization of the well-studied Γ_D -function known from ADFs. This allows us to define the main semantics of cADFs in terms of (post-)fixpoints of the Γ_Π -functions, just like in the case of ADFs.

With our generalized Γ_Π -function at hand, we can now define the main semantics for cADFs as straightforward generalizations of the ADF-semantics:

Definition 7. Let a cADF Π over At and an interpretation $v \in \mathcal{V}^4(\text{At})$ be given, then:

- v is *admissible* for Π iff there is some $v' \in \Gamma_\Pi(v)$ s.t. $v \leq_i v'$.
- v is *complete* for Π iff $v \in \Gamma_\Pi(v)$.
- v is *preferred* for Π if it is a \leq_i -maximal among all admissible interpretation for Π ;
- v is *grounded* for Π if it is a \leq_i -minimal among all complete interpretation for Π ;
- v is a *two-valued model* for Π iff $v \in \mathcal{V}^2(\text{At})$ and v is complete.

Example 8 (Example 7 ctd.). We see that for Π from Example 7, there are two complete interpretations: TF and UT. This can be seen by observing that $TF \in \Gamma_\Pi(TF)$ and $UT \in \Gamma_\Pi(UT)$. Since these interpretations are \leq_i -incomparable, both interpretations are also grounded. The admissible interpretations are: UU, UT, TU and TF. Thus, UT and TF are also preferred.

Example 9. Let $\Pi = \{b \triangleleft p, f \triangleleft b, \neg f \triangleleft p\}$ formulated over $\Sigma = \{b, f, p\}$ be given. $v_U = UUU$ is the unique complete interpretation and thus also grounded. It is also the unique admissible interpretation.

Notice that e.g. TIT is *not* complete, since $\Gamma_\Pi(TIT) = \{TIU\}$. The reason for $\Gamma_\Pi(TIT)(p) = U$ is since there is no acceptance pair $p \triangleleft \phi \in \Pi$. The intuition is that p is only accepted if we have good information to do so, but no such information is given by any $\phi \triangleleft \psi \in \Pi$.

It is interesting to note that for $\Pi' = \Pi \cup \{p \triangleleft p\}$, $TIT \in \Gamma_{\Pi'}(TIT) = \{TIU, TIT, TIF\}$.

As can be seen in the example above, if an atom a occurs in no statement of ϕ of any acceptance pair $\phi \triangleleft \psi \in \Pi$, then $v(a) = \text{U}$ for any admissible or complete interpretation v . However, should this be undesired, one can simply add the acceptance pair $a \triangleleft a$ for such an atom.

Design Choices in Γ_Π and Comparison with Γ_D

We now discuss the design choices made in generalizing the Γ -function from ADFs to cADFs.

A first generalization is caused by the fact that statements ϕ of acceptance pairs $\phi \triangleleft \psi$ are possibly non-atomic formulas. Since Γ_Π contains all interpretations v' that align, for any $\phi \triangleleft \psi \in \Pi$, the truth value of ϕ with $v(\psi)$, there might now be more than one interpretation v' which achieves this. As a case in point, consider the cADF $\Pi = \{p \vee q \triangleleft \top\}$, where acceptance of $p \vee q$ (which is required by any $v \in \mathcal{V}^4$, since $v(\top) = \text{T}$ for any $v \in \mathcal{V}^4$) can be guaranteed by any interpretation that satisfies p or q . Therefore, the Γ -function might contain multiple interpretations which all do an equally good job of aligning the truth values of statements ϕ with their respective conditions ψ . Thus, Γ_Π is defined as a *non-deterministic operator* (Pelov and Truszczyński 2004; Heyninck and Arieli 2021), in the sense that a single interpretation v might give rise to a non-singleton set of interpretations $\{v_1, \dots, v_n\} = \Gamma_\Pi(v)$. In the example above, we have e.g. $\Gamma_\Pi(v) = \{\text{TU}, \text{UT}\}$ for any $v \in \mathcal{V}^4(\{p, q\})$.

A second generalization w.r.t. the Γ -function for ADFs is the fact that alignment of statements ϕ with their corresponding condition ψ cannot always be done in an exact way. In more detail, for ADFs D , alignment by Γ_D of s is always exact, in the sense that $\Gamma_D(v)(s)$ coincides with the truth value assigned by $\prod_i [v]^2(C_s)$. This is not always possible for cADFs, since we might have conflicting specifications in a cADF. Take for example the cADF $\Pi = \{p \triangleleft \top; \neg p \triangleleft \top\}$. Clearly, for any $v \in \mathcal{V}^4(\text{At})$, there exists no $v' \in \mathcal{V}^4(\text{At})$ s.t. $v'(\phi) = v(\psi)$ for every $\phi \triangleleft \psi$. Indeed, this is one of the reasons we had to move to a four-valued logic, since now we can at least specify an interpretation v' which brings $v'(p)$ and $v'(\neg p)$ in alignment with $v(\top)$, in the sense that $v'(p)$ and $v'(\neg p)$ are at least as informative as $v(\top)$, i.e. $v'(p) \geq_i v(\top)$ and $v'(\neg p) \geq_i v(\top)$ (for any $v \in \mathcal{V}^4(\text{At})$).

Semantical Properties of cADF-semantics

In this section, we show central semantical results on the semantics of cADFs. In particular, we show some relationships between the semantics, and we show under which conditions admissible, complete, grounded and preferred interpretations are guaranteed to exist.

We start by observing that, just like for ADFs, complete interpretations are admissible:

Proposition 2. Let a cADF Π and a complete interpretation v for Π be given. Then v is admissible.

For showing the existence of all semantics, it will be useful to limit attention to what we will call *well-formed cADFs*. The main idea is that we want to avoid cADFs Π for which $\Gamma_\Pi(v) = \emptyset$ for some $v \in \mathcal{V}^4(\text{At})$, as occurs in e.g. the following example:

Example 10. $\Pi = \{p \triangleleft \top, \neg p \triangleleft \top, p \vee \neg p \triangleleft p\}$.

| | | | |
|------------|-----------------|------------|-----------------|
| v | $\Gamma_\Pi(v)$ | v | $\Gamma_\Pi(v)$ |
| T | $\{\text{I}\}$ | F | $\{\text{I}\}$ |
| U | $\{\text{I}\}$ | I | \emptyset |

Notice that $\Gamma(\text{I}) = \emptyset$.

Definition 8. A *well-formed cADF* is a cADF Π s.t. $\Gamma_\Pi(v) \neq \emptyset$ for any $v \in \mathcal{V}^4(\text{At})$.

As a side note, we observe that a syntactic sufficient condition for well-formedness of a cADF Π is to simply require that for every acceptance pair $\phi \triangleleft \psi \in \Pi$, the statement ϕ is a logically contingent formula. We call such cADFs *unconstrained*:

Definition 9. A cADF Π is *unconstrained* iff for every $\phi \triangleleft \psi \in \Pi$, ϕ is logically contingent.

We explain the term of *unconstrained cADF* as follows. Notice that an acceptance pair $\phi \triangleleft \psi$, where ϕ is a tautology or a falsity, can be seen as a constraint, in the sense that it forces ψ to be set to the value of ϕ (i.e. $v(\psi) = \text{T}$ if ψ is a tautology and $v(\psi) = \text{F}$ if ψ is a falsity) for any complete extension. To see this, observe that $v(\phi) = \text{T}[\text{F}]$ for any $v \in \mathcal{V}^4$ if ϕ is a tautology[falsity]. In particular, for any $v' \in \Gamma_\Pi(v)$, it will hold that $v(\phi) = \text{T}[\text{F}]$. It is quite interesting that the framework naturally allows for the formulation of constraints, but for the development of the meta-theory, it will prove useful to restrict attention to well-formed cADFs. It is an interesting question for future work to see whether *constrained argumentation frameworks* (Coste-Marquis, Devred, and Marquis 2006) can be captured using such constraints.

Proposition 3. Any unconstrained cADF Π is well-formed.

However, not all well-formed cADFs are unconstrained:

Example 11. Consider $\Pi = \{a \vee \neg a \triangleleft a \vee \neg a\}$. Then clearly, for any $v \in \mathcal{V}^4(\text{At})$, $\Gamma_\Pi(v) = \{\text{T}\}$ (since $[\text{U}]^2 = \{\text{T}, \text{F}\}$ and $\prod_i \{\text{T}, \text{F}\}(a \vee \neg a) = \text{T}$).

All semantics exist for well-formed cADFs:

Proposition 4. For any well-formed cADF, there exists an admissible, preferred, complete and grounded interpretation.

6 Grounded Interpretations and the Grounded State

One of the crucial properties of ADFs is that a unique grounded interpretation is guaranteed to exist. This property does not generalize to the grounded semantics of cADFs, in view of the indeterminism that cADFs allow to express. As a case in point consider $\Pi = \{p \vee q \triangleleft \top\}$, which has two \leq_i -minimal complete interpretations: v_1 and v_2 with: $v_1(p) = \text{T}$, $v_1(q) = \text{U}$, $v_2(p) = \text{U}$ and $v_2(q) = \text{T}$. Thus, there might be cADFs that do not have a unique grounded interpretation. This might be seen as problematic, since the grounded interpretation for ADFs can be calculated efficiently and straightforwardly by iterating Γ_D starting from v_U . Since the grounded interpretation v_g is \leq_i -minimally complete and unique for ADFs, it approximates any other complete interpretation of the ADF in question (in the sense

that $v_g \leq_i v$ for any complete interpretation v). We are now interested in defining a similar concept for cADFs, that is, a unique representation of the \leq_i -minimal information expressed by a cADF that can be unambiguously obtained by application of Γ_Π and approximates any complete interpretation. This can be done by looking at a set of interpretations instead of a single interpretation. We note that this idea is not new. For example, many well-founded semantics for disjunctive logic programming take up this idea, resulting in a well-founded state (Baral, Lobo, and Minker 1992; Brass and Dix 1995; Alcântara, Damásio, and Pereira 2005).² Accordingly, we will be interested in a *grounded state* $\mathcal{V}' \subseteq \mathcal{V}^4(\text{At})$ that represents the minimal knowledge entailed by a cADF. This grounded state can be defined as the \leq_i^S -minimal fixpoint of Γ'_Π , a generalization of Γ_Π to sets of interpretations. Γ'_Π is obtained as follows:

Definition 10. Given a cADF Π and $\mathcal{V}' \subseteq \mathcal{V}^4(\text{At})$:

$$\Gamma'_\Pi(\mathcal{V}') = \min_{\leq_i} \bigcup_{v \in \mathcal{V}'} \Gamma_\Pi(v)$$

We lift \leq_i to sets of interpretations by defining, for $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathcal{V}^4(\text{At})$, $\mathcal{V}_1 \leq_i^S \mathcal{V}_2$ iff for every $v_2 \in \mathcal{V}_2$ there is some $v_1 \in \mathcal{V}_1$ s.t. $v_1 \leq_i v_2$.

Definition 11. Let a cADF Π be given. $\mathcal{V}' \subseteq \mathcal{V}^4(\text{At})$ is: (1) a *complete state* (for Π) iff $\mathcal{V}' = \Gamma'_\Pi(\mathcal{V}')$, (2) a *grounded state* (for Π) iff \mathcal{V}' is a \leq_i^S -minimally complete state (for Π).

Proposition 5. Let a cADF Π be given. Then:

1. There exists a unique grounded state which can be obtained by iterating Γ'_Π , starting with v_U .
2. For any ADF D , the grounded state coincides with $\{v\}$, where v is the grounded interpretation of D .
3. Where \mathcal{V}' is the grounded state for Π and v is a complete interpretation of Π , we have that: $\mathcal{V}' \leq_i^S \{v\}$.

Example 12. Let $\Pi = \{p \vee q \triangleleft \top, s \triangleleft p, s \triangleleft q\}$ over the signature $\{p, q, s\}$. Then we can obtain the grounded state for Π by the following calculation: (1) $\Gamma'_\Pi(\{v_U\}) = \{\text{TUU}, \text{UTU}\}$; (2) $\Gamma'_\Pi(\Gamma'_\Pi(\{v_U\})) = \min_{\leq_i} (\Gamma_\Pi(\text{TUU}) \cup \Gamma_\Pi(\text{UTU})) = \{\text{TUT}, \text{UTT}\}$; (3) $\Gamma'_\Pi(\Gamma'_\Pi(\Gamma'_\Pi(\{v_U\}))) = \min_{\leq_i} (\Gamma_\Pi(\text{TUT}) \cup \Gamma_\Pi(\text{UTT})) = \{\text{TUT}, \text{UTT}\}$.

Since in the third step a fixed point was reached, we see that the grounded state of Π is $\{\text{TUT}, \text{UTT}\}$. The grounded state consists of two interpretations, which both make s true, and either make p or q true.

Remark 2. All semantics defined in this paper have been implemented in Java using the `Tweety`-library (Thimm 2017).

7 Related Work

To the best of our knowledge, no generalizations of ADFs as we have suggested here have been proposed before. As a

²Some semantics explicitly use the idea of a set of interpretations (Alcântara, Damásio, and Pereira 2005), whereas other semantics are phrased syntactically, as a set of disjunctions (Baral, Lobo, and Minker 1992; Brass and Dix 1995), which is equivalent to a set of interpretations (Seipel, Minker, and Ruiz 1997).

side effect of the semantics of cADFs, we obtain also a four-valued semantics of ADFs and argumentation frameworks. However, *epistemic graphs* (Hunter, Polberg, and Thimm 2020) can be regarded as an orthogonal approach to extend the expressivity of ADFs. There, general propositional formulas are interpreted through a probabilistic semantics (that is not related to ADF semantics), thus yielding an expressive probabilistic and argumentative formalism. Instead, we have a purely qualitative formalism that generalises the original ADF semantics directly. Attacks on sets of arguments are possible in SETAFs (Nielsen and Parsons 2006). However, in SETAFs, only attacks from sets of arguments are allowed, and not on sets of arguments. Furthermore, support is not studied in SETAFs. In future work, we will study how to formulate cADFs that model attack and support between sets of arguments. Four-valued semantics for the more specific abstract argumentation frameworks have been proposed in (Baroni, Giacomin, and Liao 2015; Arieli 2012).

The semantics of 4CM bears similarities to those of *generalized possibilistic logic* (in short, GPL) (Dubois 2012), where a pair of sets of possible worlds is used to represent the information given by a four-valued interpretation. $[v]^4$ might consist of more than two sets of worlds, which results in e.g. $v_1(p) = v_1(q) = v_1(\neg p \vee \neg q) = \text{I}$, different from GPL.

ADFs have been generalized in other works, in particular as to allow for the handling of weights (Brewka et al. 2018; Bogaerts 2019). In (Brewka et al. 2018) an instantiation of weighted ADFs using Belnap’s four-valued logic is discussed. However, in the setting of (Brewka et al. 2018) this results in five truth-values, since in weighted ADFs, the truth-values are always supplemented with an information-theoretic minimum \cup that is not part of the original set of truth-values. This is counter-intuitive, as Belnap’s truth-values already include a truth-value expressing undecidability. Furthermore, this instantiation uses Belnap’s four-valued logic to evaluate complex formulas, which means that tautologies can be both assigned Belnap’s inconsistent and incomplete truth-values (but never the external \cup -value). Finally, syntactically, weighted ADFs conform with ADFs in the sense that they require exactly one acceptance condition to be assigned to every node, and thus, the syntax of cADFs also generalizes the syntax of weighted ADFs.

8 Conclusion

In this paper, we have defined and studied cADFs, which generalize ADFs and allow for indeterminism, over- and underspecifications. Semantics for cADFs are defined in terms of a Γ -function mapping four-valued interpretations to sets of four-valued interpretations. There remains still a lot of work to be done on cADFs. As a first next step, there are still some semantics that need to be generalized from ADFs to cADFs, in particular the stable semantics. Thereafter, we plan to study the computational complexity and realizability (in the style of (Pührer 2020)) of cADFs. On the basis of these steps, we will then have a clear view of which formalisms can be captured by cADFs, e.g. disjunctive and propositional logic programming (Minker and Seipel 2002; Ferraris 2005) and logics for nonmonotonic conditionals (Kraus, Lehmann, and Magidor 1990).

Acknowledgements

The research reported here was supported by the Deutsche Forschungsgemeinschaft under grant KE 1413/11-1.

References

- Alcântara, J.; Damásio, C. V.; and Pereira, L. M. 2005. A well-founded semantics with disjunction. In *International Conference on Logic Programming*, 341–355. Springer.
- Arieli, O. 2012. Conflict-tolerant semantics for argumentation frameworks. In *European Workshop on Logics in Artificial Intelligence*, 28–40. Springer.
- Baral, C.; Lobo, J.; and Minker, J. 1992. Generalized disjunctive well-founded semantics for logic programs. *Annals of Mathematics and Artificial Intelligence*, 5(2): 89–131.
- Baroni, P.; Giacomin, M.; and Liao, B. 2015. I don’t care, I don’t know. . . I know too much! On Incompleteness and Undecidedness in Abstract Argumentation. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation*, 265–280. Springer.
- Belnap, N. D. 1977. A useful four-valued logic. In *Modern uses of multiple-valued logic*, 5–37. Springer.
- Belnap, N. D. 2019. How a computer should think. In *New Essays on Belnap-Dunn Logic*, 35–53. Springer.
- Bogaerts, B. 2019. Weighted abstract dialectical frameworks through the lens of approximation fixpoint theory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2686–2693.
- Brass, S.; and Dix, J. 1995. Disjunctive Semantics based upon Partial and Bottom-Up Evaluation. In *ICLP*, 199–213.
- Brewka, G.; Strass, H.; Ellmauthaler, S.; Wallner, J. P.; and Woltran, S. 2013. Abstract dialectical frameworks revisited. In *ICJAI*.
- Brewka, G.; Strass, H.; Wallner, J. P.; and Woltran, S. 2018. Weighted abstract dialectical frameworks. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Cayrol, C.; and Lagasque-Schiex, M.-C. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 378–389. Springer.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2006. Constrained Argumentation Frameworks. *KR*, 6: 112–122.
- Dubois, D. 2012. Reasoning about ignorance and contradiction: many-valued logics versus epistemic logic. *Soft Computing*, 16(11): 1817–1831.
- Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *AI*, 77: 321–358.
- Ferraris, P. 2005. Answer sets for propositional theories. In *International Conference on Logic Programming and Nonmonotonic Reasoning*, 119–131. Springer.
- Heyninck, J.; and Arieli, O. 2021. Approximation Fixpoint Theory for Non-Deterministic Operators and Its Application in Disjunctive Logic Programming. In *accepted for the International Conference on Principles of Knowledge Representation and Reasoning*.
- Heyninck, J.; Kern-Isberner, G.; Skiba, K.; and Thimm, M. 2019. Interpreting Conditionals in Argumentative Environments. In *NMR 2020 Workshop Notes*, 73.
- Hunter, A.; Polberg, S.; and Thimm, M. 2020. Epistemic Graphs for Representing and Reasoning with Positive and Negative Influences of Arguments. *Artificial Intelligence*, 281: 103236.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *AI*, 44(1-2): 167–207.
- Minker, J.; and Seipel, D. 2002. Disjunctive logic programming: A survey and assessment. In *Computational logic: logic programming and beyond*, 472–511. Springer.
- Nielsen, S. H.; and Parsons, S. 2006. A generalization of Dung’s abstract framework for argumentation: Arguing with sets of attacking arguments. In *International Workshop on Argumentation in Multi-Agent Systems*, 54–73. Springer.
- Pelov, N.; and Truszczynski, M. 2004. Semantics of disjunctive programs with monotone aggregates: an operator-based approach. In *Proceedings of NMR’04*, 327–334.
- Polberg, S. 2016. Understanding the abstract dialectical framework. In *European Conference on Logics in Artificial Intelligence*, 430–446. Springer.
- Pührer, J. 2020. Realizability of three-valued semantics for abstract dialectical frameworks. *Artificial Intelligence*, 278: 103198.
- Seipel, D.; Minker, J.; and Ruiz, C. 1997. Model generation and state generation for disjunctive logic programs. *The Journal of Logic Programming*, 32(1): 49–69.
- Thimm, M. 2017. The Tweety library collection for logical aspects of artificial intelligence and knowledge representation. *KI-Künstliche Intelligenz*, 31(1): 93–97.