

A Nonmonotonic Proof Theory for Dialectical Argumentation Under Bounded Resources

Kees VAN BERKEL ^{a,1} and Sanjay MODGIL ^b

^a*Institute for Logic and Computation, TU Wien, Vienna, Austria*

^b*Department of Informatics, King's College, London, UK*

Abstract. This paper makes a proof-theoretic contribution to resource-bounded dialectical argumentation. Practical deployment of argumentation-based nonmonotonic reasoning can benefit from integration of proof-theoretic means for construction and evaluation of arguments, while accommodating agents with bounded resources. We present a nonmonotonic proof system that implements a generalization of dialectical argumentation, adopting arguments that differentiate between committed and supposed premises, while integrating rules for constructing arguments. The proof system adopts annotations to capture the changing status of arguments in a derivation and employs annotation revision rules that evaluate the *dialectical* acceptability of these arguments, yielding rational outcomes under resource bounds. Soundness and completeness is shown for the dialectical grounded semantics.

Keywords. dialectical argumentation, logical argumentation, resource-bounded reasoning, proof theory, nonmonotonic logic

1. Introduction

Dung's seminal theory of argumentation [1] formalises nonmonotonic (*nm*) reasoning in terms of the exchange of arguments. Arguments are constructed via application of a base logic L 's inference rules to a belief base \mathcal{B} of premises, and then related by defeats in a Dung argumentation framework (*AF*). The *AF*'s justified arguments are evaluated under various semantics, and their claims can then be shown to equate with various *nm* consequence relations defined directly over \mathcal{B} . Argument game proof theories (e.g., [2]) can be deployed to determine whether a given argument is justified (i.e., whether the argument's conclusion is an *nm* consequence from \mathcal{B}). Indeed, these game proof theories can be generalised to dialogical formalisations of *nm* reasoning in which agents exchange arguments so as to evaluate whether an argument is justified (e.g., [3,4]). Common to these developments of Dung's theory is the assumption that *all* arguments are constructed and then instantiate an *AF* prior to their evaluation. This assumption is unlikely to reflect the practical demands of argumentation-based *nm* reasoning. Rather, one would expect an interleaving of argument construction and evaluation. For example, given an argument X whose status is to be evaluated, one would *only then* expect an agent to be motivated

¹Corresponding Author: Kees van Berkel, kees@logic.at.

to construct possible defeaters, *subject to limited (computational/cognitive/temporal) resources available* for constructing arguments in the base logic L . In fact, whereas reasoning by non-ideal, finite, resource-bounded (human or artificial) agents is well-studied in philosophy, economics, and cognitive (neuro)science, progress on a fundamentally *logical account* of bounded rationality, in the context of nonmonotonic reasoning and argumentation more generally [5], has only recently been addressed [6,7].

Arieli et al [8,9] propose a sequent-style proof-theoretic formalism that allows interleaving argument construction and evaluation. It does so by integrating a sequent calculus LC of a basic logic L , with rules that proof-theoretically establish the acceptability of LC-derivable arguments (i.e., sequents) through *annotating* sequents with their (changing) status in a derivation. However, [8,9] do not stipulate that the premises of arguments are checked for consistency or subset minimality (i.e., that no proper subset of a constructed argument's premises entails the conclusion). As shown in [6], for classical logic arguments (i.e., where L is classical logic) this may result in violation of the non-contamination rationality postulates [10]. Also, agents are tacitly assumed to have unbounded resources, in the sense that one assumes availability of *all* L -based arguments that can be constructed from \mathcal{B} (i.e., 'logical omniscience' is assumed). For a classical base logic, this may lead to violation of the consistency rationality postulates [11].

The main contribution of this paper is to further develop the annotated sequent formalism of [8] to accommodate resource-bounded agents that are not logically omniscient, and that need not check premise consistency or subset minimality (and do not thereby incur the additional significant computational resources needed for these checks), while still guaranteeing satisfaction of the consistency and non-contamination postulates.

We do so by proposing a sequent system that integrates the rules for a base logic L with rules for evaluating *dialectical* acceptability. That is, we integrate and generalize d'Agostino and Modgil's 'dialectical' approach to classical logic argumentation [6,12], wherein acceptability of arguments, and hence the extensions under various semantics, is evaluated through the deployment of *dialectical* classical logic arguments. Dialectical arguments make an epistemic distinction between premises that an agent commits to, and premises that can be supposed for the *sake of argument* (e.g., those committed by an interlocutor who may be imaginary in the case of single-agent reasoning). Dialectical arguments are *not* checked for premise consistency or subset minimality, and only minimal assumptions are made as to the resources available for constructing arguments. Yet [6] show that *all* rationality postulates are satisfied. However, unlike in Dung's theory [1], the defeaters of a dialectical argument x acceptable w.r.t. some extension \mathcal{E} , may increase in number as the set \mathcal{E} expands. This implies that we cannot straightforwardly adapt the proof-theoretic mechanism developed in [8,9] for evaluating acceptability. In this paper, we outline a proof-theoretic procedure that ensures preservation of the acceptability of dialectical arguments in an extending derivation. In particular, we show that our proof system is sound and complete for the grounded semantics of dialectical argumentation.

Section 2 recalls Dialectical Argumentation (*D-Arg*), and generalises the approach to accommodate a large class of resource-bounded base logics. In Section 3, we define an Annotated Dialectical Argumentation Calculus (ADAC), that enables interleaving of argument construction and evaluation of *dialectical* acceptability. All rationality postulates remain satisfied. In Section 4, we propose an epistemic closure procedure that preserves acceptability of arguments as an ADAC-derivation extends, and prove that ADAC is sound and complete for *D-Arg* under grounded semantics. Section 5 concludes.

2. Preliminaries: Dialectical Argumentation

We assume a monotonic consequence relation $\vdash_{\mathcal{L}}$ defined over a propositional or first order language \mathcal{L} . We use upper and lower case Greek letters as metavariables, respectively ranging over individual/sets of formulas in \mathcal{L} . We write $\Delta \parallel \Gamma$ to denote that $\Delta \subseteq \mathcal{L}$ is syntactically disjoint from $\Gamma \subseteq \mathcal{L}$ (see [6, p.20] for an exact definition). We assume that \mathcal{L} and $\vdash_{\mathcal{L}}$ satisfy the following: i) $\perp \in \mathcal{L}$; ii) $\{\varphi\} \vdash_{\mathcal{L}} \varphi$; iii) if $\Delta \cup \Gamma \vdash_{\mathcal{L}} \alpha$ and $\Delta \parallel (\Gamma \cup \{\alpha\})$ then either $\Delta \vdash_{\mathcal{L}} \perp$ (in the case that *ex falso quodlibet* holds for $\vdash_{\mathcal{L}}$) or $\Gamma \vdash_{\mathcal{L}} \alpha$; iv) for all $\varphi \in \mathcal{L} \setminus \perp$, $\bar{\varphi} \in \mathcal{L} \setminus \perp$ denotes the ‘contradictory’ of φ , and $\{\varphi, \bar{\varphi}\} \vdash_{\mathcal{L}} \perp$. For example, if $\vdash_{\mathcal{L}}$ is the classical consequence relation \vdash_{CL} , then if φ is of the form $\neg\alpha$, $\bar{\varphi} = \alpha$, else $\bar{\varphi} = \neg\varphi$. In what follows, we say $A = (\Delta, \alpha)$ is a (classical) argument whenever $\Delta \vdash_{\mathcal{L}} \alpha$, using uppercase Roman letters to denote such arguments. We write $\text{Conc}(A) = \alpha$ and $\text{Prem}(A) = \Delta$ to denote A ’s conclusion, respectively premises.

Definition 1. A dialectical argumentation framework $DF_{\mathcal{B}}$ based on $\mathcal{B} \subseteq \mathcal{L}$ is a tuple $(\mathcal{A}, \mathcal{C})$, where $\mathcal{A} \subseteq \{(\Delta, \alpha) \mid \Delta \subseteq \mathcal{B}, \Delta \vdash_{\mathcal{L}} \alpha\}$ is any set of arguments satisfying:

- P1 $\alpha \in \mathcal{B}$ implies $(\{\alpha\}, \alpha) \in \mathcal{A}$.
- P2 (Δ, α) and $(\Gamma, \bar{\alpha}) \in \mathcal{A}$ implies $(\Delta \cup \Gamma, \perp) \in \mathcal{A}$.
- P3 $(\Delta \cup \Gamma, \alpha) \in \mathcal{A}$ and $\Delta \parallel (\Gamma \cup \{\alpha\})$, implies $(\Delta, \perp) \in \mathcal{A}$ or $(\Gamma, \alpha) \in \mathcal{A}$.

$\mathcal{C} = \{(X, Y) \mid X = (\Delta, \varphi), Y = (\Pi, \psi) \in \mathcal{A}, \bar{\varphi} \in \text{Prem}(Y) = \Pi\}$ is an attack relation. We assume a strict partial ordering \prec over \mathcal{A} such that for all $(\Delta, \perp) \in \mathcal{A}$: there is an $\alpha \in \Delta$ such that $(\Delta, \perp) \not\prec (\{\alpha\}, \alpha)$ (i.e., \prec is ‘dialectically coherent’ [6]²).

We note that for $DF_{\mathcal{B}}$ s one need not incur the computational demands of checking that arguments’ premises are subset minimal or consistent, and only assume that resources suffice to satisfy P1, P2, and P3.

Remark 1. Henceforth, we assume a logic \mathbb{L} satisfying (i)-(iv). Proof of the rationality postulates in [6] relies only on \mathbb{L} satisfying these properties. Note that one may assume a resource-bounded approximation \mathbb{L}_r of \mathbb{L} , provided P1-P3 (and so implicitly (i)-(iv)) remain satisfied by the \mathbb{L}_r arguments included in a DF (e.g., see the resource-bounded approximation \vdash_r of \vdash_{CL} in [12]’s dialectical formalisation of Preferred Subtheories).

In real-world reasoning and decision-making, when evaluating the acceptability of arguments one must *epistemically* distinguish between the committed and supposed premises. To model this, *epistemic variants* of classical arguments are introduced. We use lowercase italicised Roman letters to refer to these epistemic variants.

Definition 2. Let $X = (\Delta, \alpha)$. We say $x = (\Sigma, \Gamma, \alpha)$ is an epistemic variant of X whenever $\text{Prem}(X) = \Delta = \Sigma \cup \Gamma$ and $\Sigma \cap \Gamma = \emptyset$. We also say that x is a ‘dialectical argument’. We write $\text{Com}(a) = \Delta$, $\text{Sup}(a) = \Gamma$, and $\text{Conc}(a) = \alpha$ to denote a ’s committed premises, supposed premises, and conclusion.

- Let E be any set of arguments, we write $\|X\|$ to denote the set of all epistemic variants of $X \in E$, and let $\|E\|$ denote $\bigcup_{X \in E} \|X\|$.
- Let \mathcal{E} be any set of dialectical arguments, $\text{Com}(\mathcal{E})$, $\text{Sup}(\mathcal{E})$, and $\text{Conc}(\mathcal{E})$ respectively denote $\bigcup_{x \in \mathcal{E}} \text{Com}(x)$, $\bigcup_{x \in \mathcal{E}} \text{Sup}(x)$, and $\bigcup_{x \in \mathcal{E}} \text{Conc}(x)$.

²For all $\alpha \in \Delta$: $(\Delta, \perp) \prec (\{\alpha\}, \alpha)$ means irrationally preferring to commit to all α in the inconsistent Δ .

When establishing whether $x = (\Sigma, \Pi, \alpha)$ is acceptable w.r.t. a set \mathcal{E} of dialectical arguments, it is only the committed premises Σ that can be targeted. Furthermore, an attack by $y = (\Delta, \Gamma, \alpha)$ on $\beta \in \Sigma$ is contingent on the suppositions Γ of y being commitments in x together with \mathcal{E} , i.e., $\Gamma \subseteq \text{Com}(\mathcal{E} \cup \{x\})$. The intuitive idea behind this is as follows:

“Given that I commit to Δ and supposing for the sake of argument your commitments Γ in \mathcal{E} and x , I can construct an argument y that challenges your premise $\beta \in \Sigma$.”

Such an attack succeeds as a defeat only if the targeted premise is not strictly preferred over its attacker, that is, if $Y = (\Delta \cup \Gamma, \alpha) \not\prec (\{\beta\}, \beta)$ (recall that \prec is defined over \mathcal{A}).

An argument of the form $y = (\Delta, \Gamma, \perp)$ can challenge x by arguing that the premises Γ committed in $\mathcal{E} \cup \{x\}$, together with Δ , are inconsistent. Here, x should only then be targeted if at least one of x 's committed premises β is in Γ and so is ‘culpable’ in contributing to the inconsistency. Again, y defeats x if $Y \not\prec (\{\beta\}, \beta)$. However, if $\Delta = \emptyset$ then y dialectically demonstrates a commitment to inconsistent premises Γ in $\mathcal{E} \cup \{x\}$, and so the attack succeeds as a defeat independently of preferences.

Finally, $z \in \mathcal{E}$ can defend x by defeating y , while supposing any of y 's commitments, i.e., $\text{Sup}(z) \subseteq \text{Com}(y)$. If y challenges the acceptability of x w.r.t. \mathcal{E} , it is not required that y itself be acceptable w.r.t. some set of dialectical arguments. Hence, for z 's defeat on y it suffices to only suppose the committed premises of y . Let us make the above precise.

Definition 3. Let $DF_{\mathcal{D}} = (\mathcal{A}, \mathcal{C})$ and let $\mathcal{E} \subseteq \|\mathcal{A}\|$, $y = (\Delta, \Gamma, \varphi) \in \|\mathcal{Y}\|$, $x = (\Pi, \Sigma, \psi) \in \|\mathcal{X}\|$, with $X, Y \in \mathcal{A}$:

1. if $\varphi \neq \perp$, then y defeats x w.r.t. \mathcal{E} , denoted $y \Rightarrow_{\mathcal{E}} x$, iff:
 - a) $(Y, X) \in \mathcal{C}$ on $X' = (\{\beta\}, \beta)$, with $\beta = \overline{\varphi} \in \text{Com}(x)$ and $Y \not\prec X'$;
 - b) $\Gamma \subseteq \text{Com}(\mathcal{E} \cup \{x\})$.
2. if $\varphi = \perp$, then y defeats x w.r.t. \mathcal{E} , denoted $y \Rightarrow_{\mathcal{E}} x$, iff:
 - a) $\Gamma \cap \text{Com}(x) \neq \emptyset$ and $\Gamma \subseteq \text{Com}(\mathcal{E} \cup \{x\})$;
 - b) either $\Delta = \emptyset$ or $\forall \beta \in \Gamma \cap \text{Com}(x)$, $Y \not\prec X' = (\{\beta\}, \beta)$.

In both cases we may say y defeats x “on $x' = (\{\beta\}, \emptyset, \beta)$ ” or “on β .”

Example 1. Consider the ‘foreign commitment’ example [13]. We assume a propositional \mathcal{L} , with atoms as non-italicised lower case Roman letters. Let a/b/c denote “one attends conference a/b/c” and g denote “the budget is limited.” Fig. 1 shows some of the dialectical arguments for this scenario. Argument d expresses that attending both a and b , precludes attending c . $g4$ and $g5$ express that the budget is insufficient to attend both a and b . Typically, one must unintuitively commit to a or b to challenge d (on premises b or a respectively)[13]. In D-Arg, one can suppose for the sake of argument that one attends a or b . Hence, in Fig. 1, $d \Rightarrow_{\mathcal{E}} c$ but $g4 \Rightarrow_{\emptyset} d$ on b while supposing d 's commitment to a and $g5 \Rightarrow_{\emptyset} d$ on a while supposing d 's commitment to b . The intuitive outcome is that the grounded extension \mathcal{E} justifies attending c irrespective of attending a or b .

Note that c is also defeated by the ‘explosively contaminated’ q . But by P3, resources suffice to construct f . Now, dialectical arguments that do not commit to any premises cannot (by Definition 3) be defeated, and so are acceptable w.r.t. any set of dialectical arguments. Hence $f \in \mathcal{E}$, and $f \Rightarrow_{\emptyset} q$ (on p, q and $q \rightarrow \neg p$), so defending c .

To determine the acceptability of dialectical arguments, Dung semantics [1] are adjusted to the dialectical setting [6]. Then, in defining entailment relations, only the con-

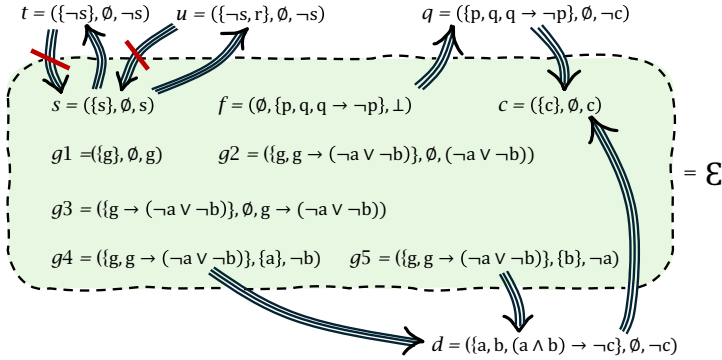


Figure 1. Dialectical arguments from Ex. 1 and 2. The encircled arguments form the set \mathcal{E} . The arrows emanating from \mathcal{E} denote \Rightarrow_{\emptyset} defeats and those towards arguments in \mathcal{E} denote $\Rightarrow_{\mathcal{E}}$ defeats. The struck through arrows denote attacks that do not succeed as defeats (given \prec).

clusions of *unconditional* arguments that commit to *all* their premises in a dialectical extension \mathcal{E} , identify the conclusions supported by \mathcal{E} . Def. 4 makes this precise.

Definition 4. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$, and $\mathcal{E} \subseteq \|\mathcal{A}\|$, $x \in \|\mathcal{A}\|$. Then:

- x is **acceptable** w.r.t. \mathcal{E} iff for all $y \in \|\mathcal{A}\|$ s.t. $y \Rightarrow_{\mathcal{E}} x$, there is a $z \in \mathcal{E}$ s.t. $z \Rightarrow_{\emptyset} y$.
- \mathcal{E} is **conflict free** iff there is no $x, y \in \mathcal{E}$ s.t. $y \Rightarrow_{\mathcal{E}} x$.

For $\text{sem} \in \{\text{admissible, grounded, preferred, stable}\}$, a dialectical sem extension is defined similarly to *Dung semantics* [1] in terms of acceptability and conflict freeness above, e.g., \mathcal{E} is dialectical admissible iff \mathcal{E} is conflict free and all $x \in \mathcal{E}$ are acceptable w.r.t. \mathcal{E} .

Let \mathcal{E} be a dialectical sem extension, then $E = \{(\Delta, \alpha) \mid (\Delta, \emptyset, \alpha) \in \mathcal{E}\}$ is a sem extension. We adopt the following sceptical (s) inference relation:

- $DF_{\mathcal{B}} \vdash_{\text{sem}}^s \alpha$ iff there is an $A \in \mathcal{A}$, $\text{Conc}(A) = \alpha$, A is in every sem extension.

As shown in [6], for \vdash_{CL} , the *consistency*, and *closure* postulates [11] are satisfied by any $DF_{\mathcal{B}}$ of Def. 3 (e.g., consistency is assured by undefeatable arguments that attack the committed inconsistent premises in an extension). Furthermore, the non-contamination postulates [10] are satisfied provided that adding syntactically disjoint premises does not strengthen arguments [6]. We emphasise that proofs of these postulates in [6] straightforwardly apply to any (resource-bounded) \vdash_{\perp} as specified in this section (Remark 1).

Example 2. If \mathcal{E} in Fig. 1 is a subset of the $DF_{\mathcal{B}}$'s dialectical grounded extension, then its grounded extension includes $G1 = (\{g\}, g)$, $G2 = (\{g \rightarrow \neg a \vee \neg b\}, g \rightarrow \neg a \vee \neg b)$, $G3 = (\{g, g \rightarrow \neg a \vee \neg b\}, \neg a \vee \neg b)$, $S = (\{s\}, s)$, $C = (\{c\}, c)$. Notice that the inconsistent q in Fig. 1 defeats c . However, the undefeatable f defeats q on its commitments, thus ensuring grounded acceptance of the dialectical argument c , hence, $DF_{\mathcal{B}} \vdash_{\text{grounded}}^s c$.

If \mathcal{B} does not contain r , and $(\{\neg s\}, \neg s) \prec (\{s\}, s)$, then $s \in \mathcal{E}$. Non-contamination is then satisfied if adding the syntactically disjoint r to \mathcal{B} does not lead to a stronger ‘redundantly contaminated’ argument u (Fig 1) that defeats s , excluding it from \mathcal{E} . For this, we require $(\{\neg s, r\}, \neg s) \prec (\{s\}, s)$ together with P3, encoding that given $U = (\{\neg s, r\}, \neg s)$, resources suffice to construct the ‘non-redundant counterpart’ $T = (\{\neg s\}, \neg s)$ of U .

3. A Dialectical Proof Calculus

Agents may have limited inferential capabilities for constructing arguments and, in practice, they are not expected to construct (and evaluate) all arguments from the outset, but rather interleave evaluation with construction. For example, when evaluating whether a is acceptable w.r.t. some \mathcal{E} , an agent may *only then* deploy proof-theoretic means to construct all arguments (given their limited inferential capabilities and/or time constraints) that defeat a . This motivates formulation of a proof theory that integrates construction of arguments and evaluation of their dialectical acceptability.

We adopt the highly modular sequent formalism (e.g., see [14] for an introduction) and develop an Annotated Dialectical Argumentation Calculus, ADAC for short. Sequent-style derivation rules [14] derive sequents of the form $\Delta \Rightarrow \Gamma$, where $\Delta \subseteq \mathcal{L}$ logically entails one of the formulas in $\Gamma \subseteq \mathcal{L}$. In what follows, we assume LC to be a sequent calculus sound and complete for the base logic L (as described in Remark 1), i.e., $\vdash_{\text{LC}} \Delta \Rightarrow \Gamma$ iff $\Delta \vdash_{\text{L}} \bigvee \Gamma$. A sequent $\Delta \Rightarrow \varphi$ straightforwardly corresponds to a classical arguments (Δ, φ) from Def. 1. Epistemic variants of these sequents can then be introduced which, following [8], are augmented with annotations that keep track of the sequents' (changing) status in a derivation. This approach yields *annotated dialectical arguments* (*adas*, for short) of the form (where s is the sequent's status in the derivation)

$$(\Pi, \Sigma) \Rightarrow^{[s]} \varphi$$

that correspond to dialectical arguments (Π, Σ, φ) (from Def. 2). We use three types of annotation [9]: $[i]$ to denote that the sequent is introduced to the derivation, $[a]$ to signify that it is finally accepted, and $[r]$ to express that it is finally rejected. We write $[*]$ when the status of the *ada* is arbitrary. Lastly, ADAC is equipped with rules that modify annotations and are used to proof-theoretically establish acceptability of dialectical arguments.

Notation 1. To enhance readability, we use the font 'x' to refer to an *ada* $(\Delta, \Gamma) \Rightarrow^{[*]} \varphi$ and use the font 'x' to refer to its corresponding dialectical argument $(\Delta, \Gamma, \varphi)$. We write $x[s]$ to denote that x has status $s \in \{i, a, r\}$. For sets of *adas* $\hat{\mathcal{D}}$ we write $\hat{\mathcal{D}}^\downarrow = \{x \mid x \in \hat{\mathcal{D}}\} \subseteq \|\mathcal{A}\|$ to denote the set of corresponding dialectical arguments. The functions $\text{Com}(\cdot)$, $\text{Sup}(\cdot)$, and $\text{Conc}(\cdot)$ are defined as usual.

The rules of ADAC are given in Fig. 2. Firstly, ADAC contains all rules of the base calculus LC, written **Rules**(LC), for deriving sequents of the form $\Delta \Rightarrow \Gamma$. The rule **AX** introduces epistemic variants of LC-sequents to a derivation and annotates them as introduced $[i]$: these are *adas*. **AX**'s side-condition ensures that the introduced *ada* qualifies as a dialectical argument according to Def. 3. We then posit rules that change the annotation of *adas* depending on the status of their defeaters. When evaluating the acceptability of an *ada* x , defeat is defined with respect to *adas* already established as acceptable in an ADAC-derivation. Relativising defeats in this way is analogous to relativisation w.r.t. some set of dialectical arguments \mathcal{E} (Def. 3). We define the set of defeaters accordingly:

Definition 5. For any x we define the set of its defeaters w.r.t a set of *adas* $\hat{\mathcal{D}}$ as:

$$\text{Def}(x, \hat{\mathcal{D}}) = \{y = (\Pi, \Sigma) \Rightarrow^{[*]} \psi \mid \vdash_{\text{LC}} \Pi \cup \Sigma \Rightarrow \psi \text{ and } y \Rightarrow_{\hat{\mathcal{D}}^\downarrow} x\}$$

where $y \Rightarrow_{\hat{\mathcal{D}}^\downarrow} x$ is as defined in Def. 3.

$$\begin{array}{c}
 \mathcal{R} \frac{\Delta_1 \Rightarrow \Gamma_1 \quad \dots \quad \Delta_n \Rightarrow \Gamma_n}{\Delta_m \Rightarrow \Gamma_m} \quad \text{for each } n\text{-premise rule } \mathcal{R} \in \mathbf{Rules}(\text{LC}) \\
 \\
 \mathbf{AX} \frac{\Delta \Rightarrow \varphi}{(\Pi, \Sigma) \Rightarrow^{[i]} \varphi} \quad \text{where } \Delta \subseteq \mathcal{B}, \Delta = \Pi \cup \Sigma, \Pi \cap \Sigma = \emptyset, \text{ and } \Delta \neq \emptyset. \\
 \\
 \mathbf{FA} \frac{x = (\Delta, \Gamma) \Rightarrow^{[i]} \varphi \quad \forall((\Pi, \Sigma) \Rightarrow^{[s]} \psi \in \text{Def}(x, \text{Acc}(\mathcal{D}))) (\Pi, \Sigma) \Rightarrow^{[r]} \psi}{x = (\Delta, \Gamma) \Rightarrow^{[a]} \varphi} \\
 \\
 \mathbf{FR} \frac{x = (\Delta, \Gamma) \Rightarrow^{[i]} \varphi \quad \exists((\Pi, \Sigma) \Rightarrow^{[s]} \psi \in \text{Def}(x, \emptyset)) (\Pi, \Sigma) \Rightarrow^{[a]} \psi}{x = (\Delta, \Gamma) \Rightarrow^{[r]} \varphi} \\
 \\
 \mathbf{EA} \frac{(\Delta, \Gamma) \Rightarrow^{[i]} \varphi}{(\Delta, \Gamma) \Rightarrow^{[a]} \varphi} \quad \text{where } \Delta \subseteq \text{Com}(\text{Acc}(\mathcal{D}))
 \end{array}$$

Figure 2. The rules of ADAC (Def. 6).

Observe that for $\text{Def}(x, \hat{\mathcal{D}})$, we identify the *adas* y whose corresponding dialectical arguments y defeat x w.r.t. the set of dialectical arguments in $\hat{\mathcal{D}}^\downarrow$, provided that (subject to satisfying P1-P3) resources suffice to derive the defeating arguments in question.

The rules **FA** and **FR** both refer to the set of defeaters in Def. 5. First, the final acceptance rule **FA** evaluates the acceptability of an *ada* relative to the set of *adas* accepted in the derivation \mathcal{D} preceding the application of **FA**, denoted by $\text{Acc}(\mathcal{D})$. Then, an application of **FA** extending \mathcal{D} , licenses derivation of an accepted $x[a]$ relative to $\text{Acc}(\mathcal{D})$ whenever its defeaters y identified by $\text{Def}(x, \text{Acc}(\mathcal{D}))$, whose suppositions reference only those commitments in x and in $\text{Acc}(\mathcal{D})$, have been rejected (i.e., $y[r] \in \mathcal{D}$).

The final rejection rule **FR** licenses derivation of a rejected $x[r]$ extending \mathcal{D} , whenever there exists a defeater y , whose suppositions reference only those commitments in x , that has been accepted (i.e., $y[a] \in \mathcal{D}$). Although **FA** and **FR** reference each other, the rules are not circularly defined since **FA** also warrants derivability of $x[a]$ when x has no defeaters (e.g., an *ada* with no commitments cannot be defeated and is thus accepted).

The epistemic acceptability rule **EA** allows one to immediately derive the acceptability of an *ada* $x = (\Delta, \Gamma) \Rightarrow^{[*]} \varphi$, all of whose committed premises are amongst the committed premises of accepted *adas* in the derivation \mathcal{D} preceding the application of **EA** (i.e., $\Delta \subseteq \text{Com}(\text{Acc}(\mathcal{D}))$). Intuitively, **EA** expresses that the epistemic variant is acceptable by virtue of the acceptability of its individual committed premises in the derivation \mathcal{D} . In the next section, we show that **EA** is justified in light of **FA** and **FR** (Prop. 1).

Let us define ADAC and ADAC-derivations (we discuss an example further below).

Definition 6. Let LC be the sound and complete sequent calculus of L , $<$ a strict partial ordering over $\{(\Delta, \varphi) \mid \vdash_{\text{LC}} \Delta \Rightarrow \varphi\}$. An Annotated Dialectical Argumentation Calculus (ADAC) contains the rules of LC, written **Rules**(LC), together with the rules **AX**, **FA**, **FR**, and **EA**. An ADAC-derivation \mathcal{D} is a finite list $\langle T_1, \dots, T_n \rangle$ of tuples

$$T_i = \langle i, \mathcal{R}, S, [j, \dots, k], \text{Acc}(\mathcal{D}_{i-1}) \rangle \quad (1 \leq i \leq n)$$

for $1 \leq i \leq n$ and where \mathcal{D}_i is the derivation up to T_i and $\mathcal{D}_0 = \emptyset$. We define T_i as follows:

- i is the index indicating the position of T_i in \mathcal{D} ;
- $\mathcal{R} \in \{\mathbf{Rules}(\text{LC}), \mathbf{AX}, \mathbf{FA}, \mathbf{FR}, \mathbf{EA}\}$ is the rule applied in T_i ;
- if $\mathcal{R} \in \mathbf{Rules}(\text{LC})$, S is an LC-sequent of the form $\Delta \Rightarrow \Gamma$;
- if $\mathcal{R} \in \{\mathbf{AX}, \mathbf{FA}, \mathbf{FR}, \mathbf{EA}\}$, S is an *ada* $\mathbf{x} = (\Delta, \Gamma) \Rightarrow^{[s]} \phi$ with status $s \in \{i, a, r\}$;
- j, \dots, k refer to the tuples T_j, \dots, T_k whose sequents serve in the conditions of the rule \mathcal{R} deriving S , with $1 \leq j, k < i$;
- $\text{Acc}(\mathcal{D}_{i-1}) = \{y \mid y[a] \in T \text{ and } T \in \mathcal{D}_{i-1}\}$.

We write $\mathbf{x} \in \mathcal{D}$ when $\mathbf{x} \in T_i$ for some $T_i \in \mathcal{D}$ and define the following entailment relation:

- $\mathcal{B} \sim_{\text{ADAC}}^{[a]} \phi$ iff there is an ADAC-derivation \mathcal{D} such that $\mathbf{x} = (\Delta, \emptyset) \Rightarrow^{[a]} \phi \in \mathcal{D}$.

Note that $\sim_{\text{ADAC}}^{[a]}$ only refers to unconditional arguments since we are only interested in conclusions drawn from the accepted committed premises (cf. Def. 4). Recall that application of **FA** and **EA** refers to the accepted *adas* in the preceding derivation and, therefore, each tuple T_i includes a reference to $\text{Acc}(\mathcal{D}_{i-1})$. Moreover, we stress that an ADAC-derivation allows for interleaving application of LC-rules that derive classical sequents, and rules **AX**, **FA**, **FR**, and **EA** that derive *adas* and establish dialectical acceptability, where the two sets of rules yield distinct sequent types as concluding sequents.

Example 3. Let $\mathcal{B} = \{r, p, q, q \rightarrow \neg p, c\}$ (and assume all arguments are equally preferred). Consider the derivation \mathcal{D} (for readability, we omit reference to tuples):

$$\begin{array}{c}
 \text{some LC-reasoning} \\
 \text{=====} \\
 \text{AX} \frac{\vdash_{\text{LC}} F = p, q, q \rightarrow \neg p \Rightarrow \perp}{\mathbf{f} = (\emptyset, \{p, q, q \rightarrow \neg p\}) \Rightarrow^{[i]} \perp} \\
 \text{FA} \frac{\mathbf{f} = (\emptyset, \{p, q, q \rightarrow \neg p\}) \Rightarrow^{[i]} \perp \quad \text{Def}(\mathbf{f}, \text{Acc}(\mathcal{D}') = \emptyset) = \emptyset}{\mathbf{f} = (\emptyset, \{p, q, q \rightarrow \neg p\}) \Rightarrow^{[a]} \perp}
 \end{array}$$

where \mathcal{D}' precedes **FA**. Since \mathbf{f} has no commitments, $\text{Def}(\mathbf{f}, \text{Acc}(\mathcal{D}')) = \emptyset$ and we derive $\mathbf{f}[a]$. Suppose we extend \mathcal{D} by deriving $c = (\{c\}, \emptyset) \Rightarrow^{[i]} c$ (recall LC satisfies P1) and we want to see whether c is acceptable. Suppose that resources suffice to derive the defeating $q = (\{r, p, q, q \rightarrow \neg p\}, \emptyset) \Rightarrow^{[i]} \neg c$ with $q \in \text{Def}(c, \text{Acc}(\mathcal{D}) = \{\mathbf{f}\})$. Since LC satisfies P3, we also derive the non-contaminated $q' = (\{p, q, q \rightarrow \neg p\}, \emptyset) \Rightarrow^{[i]} \neg c \in \text{Def}(c, \text{Acc}(\mathcal{D}))$. Clearly, \mathbf{f} defeats both. We apply **FR** twice and derive $q[r]$ and $q'[r]$. Since there are no more (uncontaminated) arguments defeating c , we extend the derivation with

$$\text{FA} \frac{c[i] \quad q[r] \quad q'[r]}{c[a]}$$

and, thus, we have $\mathcal{B} \sim_{\text{grounded}}^{[a]} c$. (We may also apply **EA** to derive $c' = (\emptyset, \{c\}) \Rightarrow^{[a]} c$.)

4. Epistemic Closure and Soundness and Completeness

Using rules akin to **FA** and **FR**, [8,9] show that through exhaustive, yet finite, rule-application, their calculi yield a set of acceptable (classical) arguments identical to the grounded extension. However, *D-Arg* does not directly warrant such an approach as new defeats may become applicable to dialectical arguments, as the sets \mathcal{E} containing these arguments expands. To illustrate, suppose $a = (\{p\}, \emptyset, p) \in \mathcal{E}$ and suppose we add

$(\{c\}, \emptyset, c)$ to \mathcal{E} . We now may have to deal with the defeater $(\{c \rightarrow \neg p\}, \{c\}, \neg p)$ which did not previously defeat a w.r.t. \mathcal{E} . For this reason, [6] impose an ‘epistemic closure’ condition that guarantees that arguments in a dialectical admissible \mathcal{E} remain acceptable when the set is expanded with newly acceptable arguments. In this section, we show how rule-based proof-theoretic reasoning with ‘epistemic closure’ likewise preserves acceptability status in ADAC-derivations, ultimately yielding a sound and complete characterisation for the grounded semantics of D -Arg.

Definition 7. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{E})$ and $\mathcal{E} \subseteq \|\mathcal{A}\|$. The set $\|\mathcal{E}\|_{max} = \{b \in \mathcal{E} \mid \text{Com}(b) \subseteq \text{Com}(\mathcal{E})\}$ contains all epistemic variants warranted by committed beliefs in \mathcal{E} . We say that \mathcal{E} is epistemically maximal whenever $\mathcal{E} = \|\mathcal{E}\|_{max}$.

We define the following epistemic closure procedure for ADAC-derivations.

Definition 8. Let \mathcal{D} be an ADAC-derivation and let $\mathcal{E} = \text{Acc}(\mathcal{D})^\downarrow$. We define an epistemic closure to be the following procedure applied to \mathcal{D} :

Step 1 for each $x \in \|\mathcal{E}\|_{max} \setminus \mathcal{E}$, introduce x if $x[*] \notin \mathcal{D}$ and apply **EA** to derive $x[a]$.

Let the resulting derivation be \mathcal{D}' with $\mathcal{E}' = \text{Acc}(\mathcal{D}')^\downarrow$, then

Step 2 for each $x \in \mathcal{E}'$ and $y \in \text{Def}(x, \mathcal{E}')$, introduce y if $y[*] \notin \mathcal{D}'$ and apply **FR** to derive $y[r]$.

Let \mathcal{D}'' be the derivation resulting from Step 1 and Step 2.

We say that a derivation \mathcal{D} is epistemically coherent if the epistemic closure procedure occurs directly after each **FA** application in \mathcal{D} .

Example 4. For illustration, consider some \mathcal{D} such that $x = (\{p\} \cup \Delta, \{c\} \cup \Gamma) \Rightarrow^{[a]} \varphi \in \mathcal{D}$, $c \notin \text{Com}(\text{Acc}(\mathcal{D}))$, and there is a $y = (\{c, r \wedge \neg r\}, \emptyset) \Rightarrow^{[r]} \neg p \in \mathcal{D}$ which was rejected by $z = (\emptyset, \{r \wedge \neg r\}) \Rightarrow^{[a]} \perp$ (cf. P2 and P3). Suppose we may apply **FA** to derive $w = (\{c\}, \emptyset) \Rightarrow^{[a]} c$. An epistemic closure follows and, i.a., $w' = (\emptyset, \{c\}) \Rightarrow^{[a]} c$ and $x' = (\{p, c\} \cup \Delta, \Gamma) \Rightarrow^{[a]} \varphi$ are derived by **EA**. Now $y' = (\{r \wedge \neg r\}, \{c\}) \Rightarrow^{[i]} \neg p$ defeats x but since $(\emptyset, \{r \wedge \neg r\}) \Rightarrow^{[a]} \perp$, the closure applies **FR**, deriving $y'[r]$: x remains acceptable.

Prop. 1 guarantees that epistemic closure is well-defined. In particular, clause (iv) states that after closure, all defeaters of each accepted argument can be legitimately rejected and, so, step 2 of the closure is well-defined. Consequently, although application of **FA** may introduce new defeats on previously accepted arguments, after epistemic closure each new defeater can be rejected and accepted arguments remain acceptable.

Proposition 1. Let \mathcal{D} be an ADAC-derivation and \mathcal{D}' result after epistemic closure:

- (i) $\text{Com}(\text{Acc}(\mathcal{D})^\downarrow) = \text{Com}(\text{Acc}(\mathcal{D}')^\downarrow)$;
- (ii) \mathcal{D}' is finite;
- (iii) $\text{Acc}(\mathcal{D}')^\downarrow = \|\text{Acc}(\mathcal{D})^\downarrow\|_{max}$;
- (iv) For $x \in \text{Acc}(\mathcal{D}')$ and $y \in \text{Def}(x, \text{Acc}(\mathcal{D}'))$, there is a $z[a] \in \mathcal{D}'$ s.t. $z \in \text{Def}(y, \emptyset)$;
- (v) \mathcal{D}' is an ADAC-derivation.

Proof. Due to space reasons we omit their proofs (see Ex. 4 for an illustration of iv). \square

Corollary 1. Let \mathcal{D} be epistemically coherent: \mathcal{D} is an ADAC-derivation, $\text{Acc}(\mathcal{D})^\downarrow = \|\text{Acc}(\mathcal{D})^\downarrow\|_{max}$, and for each $x \in \text{Acc}(\mathcal{D})$ and each $y \in \text{Def}(x, \text{Acc}(\mathcal{D}))$, $y[r] \in \mathcal{D}$.

We now prove our central result: the inference relation of ADAC proof calculi corresponds to inference under the grounded semantics of DF s. In what follows, we assume a base logic L satisfying P1-P3, a finite set \mathcal{B} , such that each $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$ satisfies P1-P3 (\mathcal{A} is not necessarily finite). We assume that ADAC and $DF_{\mathcal{B}}$ reference the same L , \mathcal{B} , and \prec . First, Prop. 2 demonstrates that accepted ADAC-derivable arguments are members of a $DF_{\mathcal{B}}$'s dialectical grounded extension. To prove this, we use the characteristic function F for dialectical argumentation, which also uses epistemic closure [6].

Definition 9. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$ and $\mathcal{E} \subseteq \|\mathcal{A}\|$. Let the characteristic function be

$$F(\mathcal{E}) = \{a \in \|\mathcal{A}\| \mid \text{for all } b \Rightarrow_{\mathcal{E}} a, \text{ there is a } c \in \mathcal{E} \text{ s.t. } c \Rightarrow_{\emptyset} b\}.$$

Let $\|F(\mathcal{E})\|_{\max}$ be the epistemic maximisation of $F(\mathcal{E})$. The least fixed-point (lfp) over $DF_{\mathcal{B}}$ is defined as $\mathcal{E}^* = \bigcup_{i < \infty} \mathcal{E}_i$, where $\mathcal{E}_0 = \emptyset$ and $\mathcal{E}_i = \|F(\mathcal{E}_{i-1})\|_{\max}$.

Recall from [6] that the lfp is the dialectical grounded extension \mathcal{E}_g of $DF_{\mathcal{B}}$.

Proposition 2. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$ and $\mathcal{E}_g \subseteq \|\mathcal{A}\|$ its dialectical grounded extension. For each \mathcal{B} -based ADAC-derivation \mathcal{D} such that $\{a \mid \mathbf{a} \in \mathcal{D}\} \subseteq \|\mathcal{A}\|$: $\text{Acc}(\mathcal{D})^{\downarrow} \subseteq \mathcal{E}_g$.

Proof. Let $\langle \mathbf{a}_1[a], \dots, \mathbf{a}_n[a] \rangle$ be the list of all $\text{Acc}(\mathcal{D}) = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ in order of their introduction in \mathcal{D} . We show by induction on the length i of this list that $a_i \in \mathcal{E}_g$. In what follows, we use \mathcal{D}_i to refer to the derivation up until the derivation of $\mathbf{a}_i[a]$.

Base case (i=1). Straightforward since \mathbf{a}_1 has no defeaters.

Inductive step (i \mapsto i+1). Assume that $a_j \in \mathcal{E}_g$ for each $j \leq i$. We show that $a_{i+1} \in \mathcal{E}_g$. Suppose towards a contradiction that $a_{i+1} \notin \mathcal{E}_g$. Since $a_{i+1} \in \|\mathcal{A}\|$, there exists a $b \in \|\mathcal{A}\|$ such that $b \Rightarrow_{\mathcal{E}_g} a_{i+1}$ and

$$\text{there is no } c \in \mathcal{E}_g \text{ such that } c \Rightarrow_{\emptyset} b. \quad (1)$$

By assumption, $\text{Sup}(b) \subseteq \text{Com}(\mathcal{E}_g \cup \{a_{i+1}\})$. By IH, $\text{Acc}(\mathcal{D}_i)^{\downarrow} \subseteq \mathcal{E}_g$. Let $\mathbf{b} = (\Delta, \Gamma) \Rightarrow \varphi$ and define $\Gamma_2 = \Gamma \cap \text{Com}(\mathcal{E}_g)$. Let b' the epistemic variant $\mathbf{b}' : (\Delta \cup \Gamma_2, \Gamma \setminus \Gamma_2) \Rightarrow^{[*]} \varphi$. Clearly, $b' \Rightarrow_{\text{Acc}(\mathcal{D}_i)^{\downarrow}} a_{i+1}$. Since $\mathbf{a}_{i+1} \in \text{Acc}(\mathcal{D})$, then by definition of **FA**, there exists a $c[a] \in \mathcal{D}_i$ such that $c \Rightarrow_{\emptyset} b'$. By IH, $c \in \mathcal{E}_g$. Let $\mathbf{c} = (\Pi, \Sigma) \Rightarrow^{[*]} \varphi$ with $\Sigma \subseteq (\Delta \cup \Gamma_2)$. Let $\mathbf{c}' : (\emptyset, \Pi \cup \Sigma) \Rightarrow^{[*]} \psi$ be an epistemic variant. Trivially, $\mathbf{c}' \in \mathcal{E}_g$ since \mathbf{c}' is undefeatable. Let $\Sigma_2 = \Sigma \cap \text{Com}(\mathcal{E}_g)$. Since $\|F(\mathcal{E}_g)\|_{\max} = \mathcal{E}_g$, we have $\mathbf{c}'' : (\Pi \cup \Sigma_2, \Sigma \setminus \Sigma_2) \Rightarrow^{[*]} \varphi$ with $\mathbf{c}'' \in \mathcal{E}_g$. Since $\Sigma \setminus \Sigma_2 \subseteq \Delta = \text{Com}(b)$, we have $\mathbf{c}'' \Rightarrow_{\emptyset} b$, contradicting (1). \square

Lem. 1 (which follows from Prop. 1-2) expresses that epistemic closure guarantees that a set of acceptable *adas* is also dialectically admissible.

Lemma 1. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$ with $\mathcal{E}_g \subseteq \|\mathcal{A}\|$ and let \mathcal{D} be an ADAC-derivation. Let \mathcal{D}' result from the epistemic closure of \mathcal{D} such that $\{a \mid \mathbf{a} \in \mathcal{D}'\} \subseteq \|\mathcal{A}\|$. Then $\text{Acc}(\mathcal{D}')^{\downarrow}$ is a dialectical admissible extension of $DF_{\mathcal{B}}$.

Prop. 3 shows the existence of an ADAC-derivation whose accepted arguments form the dialectical grounded extension of a given $DF_{\mathcal{B}}$. To prove this, we need Def. 10 stipulating finitely exhaustive $DF_{\mathcal{B}}$ s, which are finite $DF_{\mathcal{B}}$ s in which all (resources-bounded) L -derivable defeaters are present. This notion ensures applicability of the finite rules **FA** and **FR**. To see that Def. 10 is well-defined, observe that each argument has only finitely many premises and so only finitely many defeaters (exhausting the finite \mathcal{B}).

Definition 10. Let $\mathcal{A} = \{A : (\Delta, \varphi) \mid \Delta \vdash_{\perp} \varphi \text{ and } \Delta \subseteq \mathcal{B}\}$ be the set of all L-derivable arguments based on (a finite) \mathcal{B} and let $(\mathcal{A}, \mathcal{C})$ be its $DF_{\mathcal{B}}$. We call $DF'_{\mathcal{B}} = \langle \mathcal{A}', \mathcal{C}' \rangle$ finitely exhaustive when $\mathcal{A}' \subseteq \mathcal{A}$ is finite, it satisfies P1-P3, $\mathcal{C}' = (\mathcal{C} \cap (\mathcal{A}' \times \mathcal{A}'))$, and for each $A \in \mathcal{A}'$ if $(B, A) \in \mathcal{C}$, then $B \in \mathcal{A}'$.

Proposition 3. Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$ be finitely exhaustive and $\mathcal{E}_g \subseteq \|\mathcal{A}\|$ its dialectical grounded extension: There exists a \mathcal{B} -based ADAC-derivation \mathcal{D} with (i) $\text{Acc}(\mathcal{D})^{\downarrow} = \mathcal{E}_g$ and (ii) \mathcal{D} is epistemically coherent.

Proof. Ad (i), we construct \mathcal{D} in analogy to the *lfp* construction: $\mathcal{E}_g = \bigcup_{i \leq \infty} \mathcal{E}_i$ with $\mathcal{E}_0 = \emptyset$ and $\mathcal{E}_i = \|F(\mathcal{E}_{i-1})\|_{\max}$. Since $DF_{\mathcal{B}}$ is finite, the fixed point procedure is finite, and so is \mathcal{D} . We define \mathcal{D} through a step-wise procedure, making use of the following lemma:

Lemma 2 (Follows directly from Lemma 19 in [6]). Let $DF_{\mathcal{B}} = (\mathcal{A}, \mathcal{C})$, $a, b \in \|\mathcal{A}\|$, and $\mathcal{E}_i \subseteq \|\mathcal{A}\|$ be a dialectical admissible set. For each $a, b \in F(\mathcal{E}_i) \setminus \mathcal{E}_i$, we have 1) $b \in F(\|\mathcal{E}_i \cup \{a\}\|_{\max})$ and 2) $\|\mathcal{E}_i \cup \{a, b\}\|_{\max} = \|\|\mathcal{E}_i \cup \{a\}\|_{\max} \cup \{b\}\|_{\max}$.

Lem. 2 states that 1) for any a, b acceptable to \mathcal{E}_i , adding them consecutively does not change their acceptability and 2) the epistemically maximal set obtained from adding a and b simultaneously, is identical to the one obtained by adding them consecutively.

We define the following step-wise procedure:

Step 1. Let $F(\mathcal{E}_0) = \{a_1, \dots, a_n\}$. Take $a_1 \in F(\mathcal{E}_0)$. Derive and introduce \mathbf{a}_1 by means of **Rules(LC)** followed by an application of **AX**. Since $a_1 \in F(\mathcal{E}_0)$ has no defeaters, apply **FA** to derive $\mathbf{a}_1[a]$. Let \mathcal{D}' be the resulting derivation. Apply epistemic closure to \mathcal{D}' (enabled by Prop. 1) to derive $\mathbf{a}'_1[a]$ for each $a'_1 \in \|\mathbf{a}_1\|_{\max}$ and $\mathbf{b}[r]$ for each $\mathbf{b} \in \text{Def}(\mathbf{a}'_1, \{\mathbf{a}_1\})$ (in this particular case there are no defeaters). Subsequently, repeat this procedure for each $a_i \in F(\mathcal{E}_0)$ with $1 < i \leq n$. Let \mathcal{D}_1 be the resulting derivation. By Lem. 1 and Lem. 2, we know that $\text{Acc}(\mathcal{D}_1)^{\downarrow} = \mathcal{E}_1 = \|F(\mathcal{E}_0)\|_{\max}$ and by Prop. 1 and the epistemic closure, we know that for each $\mathbf{a} \in \text{Acc}(\mathcal{D}_1)$ and $\mathbf{b} \in \text{Def}(\mathbf{a}, \text{Acc}(\mathcal{D}_1))$, $\mathbf{b}[r] \in \mathcal{D}_1$ (note that by assumption also $\mathbf{b} \in \|\mathcal{A}\|$).

Step i+1. Consider \mathcal{D}_i for \mathcal{E}_i and let $F(\mathcal{E}_i) \setminus \mathcal{E}_i = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$. Consider \mathbf{a}_1 . Extend \mathcal{D}_i with the derivation and introduction of \mathbf{a}_1 (using **Rules(LC)** and **AX**). Note that $\mathbf{a}_1 \notin \mathcal{D}_i$ since only members of \mathcal{E}_i are in \mathcal{D} together with all its defeaters and, since \mathcal{E}_g is conflict-free, \mathbf{a}_1 is not in \mathcal{D}_i as a rejected sequent. Let the resulting derivation be \mathcal{D}' . Apply epistemic closure to \mathcal{D}' (as in step 1) and **FR** to all newly derived $\mathbf{b} \in \text{Def}(\mathbf{a}_1, \text{Acc}(\mathcal{D}'))$, where $\mathcal{E}_i \subseteq \text{Acc}(\mathcal{D})^{\downarrow}$, to derive $\mathbf{b}[r]$ (warranted by Prop. 1). Repeat this procedure for all $\mathbf{a}_j \in F(\mathcal{E}_i) \setminus \mathcal{E}_i$, with $1 < j \leq n$. Let \mathcal{D}_{i+1} be the resulting derivation. Similar to Step 1, by Lem. 1 and 2 we have that $\text{Acc}(\mathcal{D}_{i+1})^{\downarrow} = \mathcal{E}_{i+1} = \|F(\mathcal{E}_i)\|_{\max}$ and by Prop. 1 and epistemic closure, we have that for each $\mathbf{a} \in \text{Acc}(\mathcal{D}_{i+1})$ and $\mathbf{b} \in \text{Def}(\mathbf{a}, \text{Acc}(\mathcal{D}_{i+1}))$, $\mathbf{b}[r] \in \mathcal{D}_{i+1}$.

The *lfp* construction of \mathcal{E}_g is finite by assumption and, so, the above process halts at some point. Let the resulting derivation be \mathcal{D} . Clearly, by construction $\text{Acc}(\mathcal{D})^{\downarrow} = \mathcal{E}_g$.

Ad (ii), follows directly from the construction of \mathcal{D} . \square

Prop. 2 and 3 imply that ADAC is a sound and complete proof-theoretic framework for resource-bounded dialectical argumentative reasoning under grounded semantics:

Corollary 2. $\mathcal{B} \sim_{\text{ADAC}}^{\text{a}} \varphi$ if and only if $DF_{\mathcal{B}} \sim_{\text{grd}} \varphi$, where $DF_{\mathcal{B}}$ is finitely exhaustive.

5. Related and Future Work

This paper's primary contribution, is to proof-theoretic approaches to logical argumentation. For a discussion of work related to Dialectical Argumentation we refer to [6,12]. Here, we only point out that the use of suppositions in dialectical argumentation is reminiscent of hypothetical argumentation; e.g., [15]. A comparison is left for future work. We also note that the adopted final acceptability rule references the given derivation and does not depend on possible extensions of the derivation such as in [16]. Furthermore, dialectical argumentation identifies minimal conditions on the base logic L (e.g., P1-P3) that guarantee satisfaction of rationality postulates without requiring, e.g., consistency checks and omniscience.

Last, our approach extends the formalism developed by Arieli et al [8,9], who additionally show correspondence results for credulous reasoning under stable semantics in classical argumentation. The main difference with this paper is our inclusion of resource-bounded dialectical argumentation, preserving the satisfaction of various rationality postulates. Furthermore, we adopted a finite epistemic closure procedure to warrant these results. Future work will address results for dialectical stable semantics.

References

- [1] Dung PM. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*. 1995;77(2):321-57.
- [2] Modgil S, Caminada M. Proof theories and algorithms for abstract argumentation frameworks. In: *Argumentation in AI*. Springer; 2009. p. 105-29.
- [3] Fan X, Toni F. A general framework for sound assumption-based argumentation dialogues. *Artificial Intelligence*. 2014;216(0):20-54.
- [4] Prakken H. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*. 2005;15:1009-40.
- [5] Loui RP. Process and Policy: Resource-Bounded NonDemonstrative Reasoning. *Computational intelligence*. 1998;14(1):1-38.
- [6] D'Agostino M, Modgil S. Classical Logic, Argument and Dialectic. *Artificial Intelligence*. 2018;262:15-51.
- [7] D'Agostino M, Modgil S. A Fully Rational Account of Structured Argumentation Under Resource Bounds. In: *Proc. 29th Int. Joint Conference on Artificial Intelligence*; 2020. p. 1841-7.
- [8] Arieli O, van Berkel K, Straßer C. Annotated Sequent Calculi for Paraconsistent Reasoning and Their Relations to Logical Argumentation. In: *Proceedings 31st International Joint Conference on Artificial Intelligence*; 2022. p. 2532-8.
- [9] Arieli O, van Berkel K, Strasser C. Defeasible Normative Reasoning: A Proof-Theoretic Integration of Logical Argumentation. In: *Proc. 38th Annual Conference on Artificial Intelligence*; 2024. p. 10450-8.
- [10] Caminada M, Carnielli W, Dunne P. Semi-Stable Semantics. *Journal of Logic and Computation*. 2012;22(5):1207-54.
- [11] Caminada M, Amgoud L. On the evaluation of argumentation formalisms. *Artificial Intelligence*. 2007;171(5-6):286-310.
- [12] D'Agostino M, Modgil S. A Study of Argumentative Characterisations of Preferred Subtheories. In: *Proceedings 27th International Joint Conference on Artificial Intelligence*; 2018. p. 1788-94.
- [13] Caminada M, Modgil S, Oren N. Preferences and Unrestricted Rebut. In: *Proc. Computational Models of Argument*. IOS Press; 2014. p. 209-20.
- [14] Troelstra AS, Schwichtenberg H. *Basic proof theory*. 43. Cambridge University Press; 2000.
- [15] Pardo P, Straßer C. Modular orders on defaults in formal argumentation. *Journal of Logic and Computation*. 2022 Nov. Available from: <http://dx.doi.org/10.1093/logcom/exac084>.
- [16] Arieli O, Straßer C. Logical argumentation by dynamic proof systems. *Theoretical Computer Science*. 2019;781:63-91.