Carneades and Abstract Dialectical Frameworks: A Reconstruction

Gerhard Brewka [a] and Thomas F. Gordon [b]
[a] *Universität Leipzig, Augustusplatz 10-11, 04109 Leipzig, Germany*
brewka@informatik.uni-leipzig.de
[b] *Fraunhofer FOKUS, Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany*
thomas.gordon@fokus.fraunhofer.de

Abstract. Carneades is a rather general framework for argumentation. Unlike many other approaches, Carneades captures a number of aspects, like proof burdens, proof standards etc., which are of central importance, in particular in legal argumentation.

In this paper we show how Carneades argument evaluation structures can be reconstructed as abstract dialectical frameworks (ADFs), a recently proposed generalization of Dung argumentation frameworks (AFs). This not only provides at least an indirect link between Carneades and AFs, it also allows us to handle arbitrary argument cycles, thus lifting a restriction of Carneades. At the same time it provides strong evidence for the usefulness of ADFs as analytical/semantical tools in argumentation.

1. Introduction

The Carneades model of argumentation, introduced by Gordon, Prakken and Walton in [8] and developed further in a series of subsequent papers, e.g. [9], some of them appearing in this volume [2,10], is an advanced general framework for argumentation.[1] It captures both static aspects, related to the evaluation of arguments in a particular context based on proof standards for statements and on weights arguments are given by an audience, and dynamic aspects, covering for instance the shift of proof burdens in different stages of the argumentation process.

Unlike many other approaches, Carneades does not rely on Dung's argumentation frameworks (AFs) [5] for the definition of its semantics, more specifically its notion of acceptable arguments. One goal of this paper is to provide a link, albeit an indirect one, between Carneades and AFs. As we will see, both are instances of a more general framework. Moreover, in spite of this generality, Carneades suffers from a restriction: it is assumed that the graphs formed by arguments are acyclic. This is not as bad as it may first sound, as the use of pro and con arguments allows some conflicts to be represented which require cyclic representations in other frameworks, e.g. in Dung argumentation frameworks [5]. Still, cycles in argumentation appear so common that forbidding them

[1] As of June 2010, [8] is among the 10 most cited papers which appeared in the Artificial Intelligence Journal over the last 5 years.

right from the start is certainly somewhat problematic. And indeed, the authors in [8] write (page 882):

> *"We ... leave an extension to graphs that allow for cycles through exceptions for future work."*

Finishing this open task, that is, overcoming the mentioned limitation, is another main goal of this paper. We achieve this by translating Carneades argument evaluation structures to a framework which is able to handle cycles, and which offers a selection of adequate semantics. The target framework we will be using here are abstract dialectical frameworks (ADFs), and our second main goal is to provide evidence that these frameworks are indeed useful tools in argumentation.

ADFs are a powerful generalization of Dung-style argumentation frameworks [5] recently proposed by Brewka and Woltran in [4]. Dung argumentation frameworks have an implicit, fixed criterion for the acceptance of a node in the argument graph: a node is accepted iff all its parents are defeated. This acceptance criterion can be viewed as an implicit boolean function assigning a status to an argument based on the status of its parents. The basic idea underlying ADFs is to make this boolean function explicit, and then to allow arbitrary acceptance conditions for nodes to be specified.

As shown in [4], the standard semantics for Dung frameworks - grounded, preferred and stable - can be generalized to ADFs, the latter two to a slightly restricted class of ADFs called bipolar, where each link in the graph either supports or attacks its target node. Since all ADFs we are dealing with in this paper are bipolar, we will simply speak of ADFs and omit the adjective "bipolar" whenever there is no risk of confusion. Brewka and Woltran also discuss how acceptance conditions can conveniently be specified using weights of the links in an ADF. This also makes it possible to capture proof standards in a straightforward way. All this will come in handy for our reconstruction.

There are some issues that need to be addressed before we start our reconstruction. First of all, Carneades is a moving target: the framework has developed over time, and still is developing. We thus need to fix the particular version we are dealing with. We decided to choose the version presented in the book chapter [9], partly because we assume this chapter will have many readers, partly because this version is quite well-suited for our purposes, as we will see later.

Secondly, the dynamic features of Carneades have no counterpart in ADFs. ADFs were invented to capture the static evaluation of arguments, or more generally statements, given flexible forms of dependencies among them. We will restrict our discussion in this paper to the static part of Carneades. This is entirely sufficient for the purposes of this paper, and it allows us to slightly simplify the definitions from [9], stripping off dynamic aspects irrelevant to our goals. In particular, we do not discuss different argumentation stages. What we are interested in is the evaluation of stage specific Carneades argument structures.

The paper is organized as follows. We will first present Carneades, using simplified versions of the definitions in [9] which capture the relevant stage specific notions. We then present ADFs together with their semantics. The subsequent chapter contains the main results of the paper: it shows how to translate Carneades argument structures to ADFs, proves that the translation yields the desired results for acyclic Carneades structures, and discusses how this allows us to handle arbitrary cycles in argument structures.

2. Carneades

We start with the definition of arguments in Carneades [9]:

Definition 1 (argument)**.** Let $\mathcal{L}$ be a propositional language. An **argument** is a tuple $\langle P, E, c\rangle$ where $P \subset \mathcal{L}$ are its **premises**, $E \subset \mathcal{L}$ with $P \cap E = \varnothing$ are its **exceptions** and $c \in \mathcal{L}$ is its **conclusion**. For simplicity, c and all members of P and E must be literals, i.e. either an atomic proposition or a negated atomic proposition. Let p be a literal. If p is c, then the argument is an argument **pro** p. If p is the complement of c, then the argument is an argument **con** p.

An argument evaluation structure was defined in [9] as a triple consisting of a stage, an audience, and a function assigning a proof standard to propositions. Since, as mentioned in the introduction, we are only interested here in stage specific argument evaluation, we skip the status part of the definition of stages (see [9]), keeping only the set of arguments. Furthermore, an audience is a pair consisting of a set of assumptions and a weight function. For simplicity we will represent these two parts explicitly, turning the triple into a quadruple:

Definition 2 (argument evaluation structure)**.** A **(stage specific) Carneades argument evaluation structure** (CAES) is a tuple $\langle \textit{arguments}, \textit{assumptions}, \textit{weights}, \textit{standard}\rangle$, where

1. *arguments* is an acyclic[2] set of arguments,
2. *assumptions* is a consistent set of literals, those assumed by the current audience,
3. *weights* is a function assigning a real number n, $0 \le n \le 1$, to each argument, and
4. *standard* is a total function mapping propositions in $\mathcal{L}$ to a proof standard (to be defined below).

The acceptability of a proposition p in a CAES depends on its proof standard. Carneades distinguishes 5 such standards, each one based on a particular way of aggregating applicable pro and con arguments. The notion of applicability may in turn depend on the acceptability of (other) propositions:

Definition 3 (applicability)**.** Let $\mathcal{S} = \langle \textit{arguments}, \textit{assumptions}, \textit{weights}, \textit{standard}\rangle$ be a CAES. An argument $\langle P, E, c\rangle \in \textit{arguments}$ is **applicable** in $\mathcal{S}$ if and only if

- $p \in P$ implies $p \in \textit{assumptions}$ or [$\overline{p} \notin \textit{assumptions}$ and p is acceptable in $\mathcal{S}$], and
- $p \in E$ implies $p \notin \textit{assumptions}$ and [$\overline{p} \in \textit{assumptions}$ or p is not acceptable in $\mathcal{S}$].

What remains to be defined are the proof standards *scintilla of evidence*, *preponderance of evidence*, *clear and convincing evidence*, *beyond reasonable doubt* and *dialectical validity*, which we will abbreviate as se, pe, ce, bd and dv, respectively. We directly define acceptability under a particular proof standard.

Definition 4 (acceptability)**.** Let $\mathcal{S} = \langle \textit{arguments}, \textit{assumptions}, \textit{weights}, \textit{standard}\rangle$ be a CAES. A proposition $p \in \mathcal{L}$ is **acceptable** in $\mathcal{S}$ if and only if one of the following conditions holds:

---

[2] A set of arguments is acyclic if its dependency graph is. The dependency graph has a node for each propositional atom appearing in some argument. Furthermore, there is a link from q to p whenever p depends on q, that is, whenever there is an argument pro or con p with q or $\neg q$ in its set of premises or exceptions.

- *standard*(p) = se and there is at least one applicable argument for p,
- *standard*(p) = pe, p satisfies se, and the maximum weight assigned to an applicable argument pro p is greater than the maximum weight of an applicable argument con p,
- *standard*(p) = ce, p satisfies pe, and the maximum weight of applicable pro arguments exceeds some threshold α, and the difference between the maximum weight of the applicable pro arguments and the maximum weight of the applicable con arguments exceeds some threshold β,
- *standard*(p) = bd, p satisfies ce, and the maximum weight of the applicable con arguments is less than some threshold γ,
- *standard*(p) = dv, and there is at least one applicable argument pro p and no applicable argument con p.

Note that, although acceptability is defined in terms of applicability, and applicability in terms of acceptability, the definitions are well-founded. This is due to the fact that the set of arguments is not allowed to contain cycles. However, this interdependency makes it quite difficult to generalize the definitions directly to the cyclic case. We will see how ADFs can be used to overcome this limitation.

3. Abstract Dialectical Frameworks

An ADF [4] is a directed graph whose nodes represent arguments or statements which can be accepted or not. The links represent dependencies: the status of a node s only depends on the status of its parents (denoted $par(s)$), that is, the nodes with a direct link to s. In addition, each node s has an associated acceptance condition C_s specifying the conditions under which s is accepted. This is where ADFs go beyond Dung argumentation frameworks. C_s is a boolean function yielding for each assignment of values to $par(s)$ one of the values in, out for s. As usual, we will identify value assignments with the sets of nodes which are in. Thus, if for some $R \subseteq par(s)$ we have $C_s(R) = in$, then s will be accepted provided the nodes in R are accepted and those in $par(s) \setminus R$ are not accepted.

Definition 5. An *abstract dialectical framework* is a tuple $D = (S, L, C)$ where

- S is a set of statements,
- $L \subseteq S \times S$ is a set of links,
- $C = \{C_s\}_{s \in S}$ is a set of total functions $C_s : 2^{par(s)} \to \{in, out\}$, one for each statement s. C_s is called acceptance condition of s.

S and L obviously form a graph, and we sometimes refer to elements of S as nodes. For the purposes of this paper we will only deal with a subset of ADFs, called bipolar in [4]. In such ADFs each link is either attacking or supporting:

Definition 6. Let $D = (S, L, C)$ be an ADF. A link $(r, s) \in L$ is

1. *supporting* iff, for no $R \subseteq par(s)$, $C_s(R) = in$ and $C_s(R \cup \{r\}) = out$,
2. *attacking* iff, for no $R \subseteq par(s)$, $C_s(R) = out$ and $C_s(R \cup \{r\}) = in$.

For simplicity we will only speak of ADFs here, keeping in mind that all ADFs in this paper are indeed bipolar.

It turns out that Dung's standard semantics - grounded, stable, preferred - can be generalized adequately to ADFs. We first introduce the notion of a model. Intuitively, in a model all acceptance conditions are satisfied.

Definition 7. Let $D = (S, L, C)$ be an ADF. $M \subseteq S$ is a *model* of D if for all $s \in S$ we have $s \in M$ iff $C_s(M \cap par(s)) = in$.

We first define the generalization of grounded semantics:

Definition 8. Let $D = (S, L, C)$ be an ADF. Consider the operator

$$\Gamma_D(A, R) = (acc(A, R), reb(A, R))$$

where

$$acc(A, R) = \{r \in S \mid A \subseteq S' \subseteq (S \setminus R) \Rightarrow C_r(S' \cap par(r)) = in\}, \text{ and}$$
$$reb(A, R) = \{r \in S \mid A \subseteq S' \subseteq (S \setminus R) \Rightarrow C_r(S' \cap par(r)) = out\}.$$

Γ_D is monotonic in both arguments and thus has a least fixpoint. E is the *well-founded model* of D iff for some $E' \subseteq S$, (E, E') is the least fixpoint of Γ_D.

For stable models we apply a construction similar to the Gelfond/Lifschitz reduct for logic programs. The purpose of the reduction is to eliminate models in which nodes are in just because of self supporting cycles:

Definition 9. Let $D = (S, L, C)$ be an ADF. A model M of D is a *stable model* if M is the least model of the reduced ADF D^M obtained from D by

1. eliminating all nodes not contained in M together with all links in which any of these nodes appear,
2. eliminating all attacking links,
3. restricting the acceptance conditions C_s for each remaining node s to the remaining parents of s.

Preferred extensions in Dung's approach are maximal admissible sets, where an admissible set is conflict-free and defends itself against attackers. This can be rephrased as follows: E is *admissible* in a Dung argumentation framework $A = (AR, att)$ iff for some $R \subseteq AR$

- R does not attack E, and
- E is a stable extension of $(AR\text{-}R, att \cap (AR\text{-}R \times AR\text{-}R))$.

This leads to the following generalization:

Definition 10. Let $D = (S, L, C)$, $R \subseteq S$. D-R is the ADF obtained from D by

1. deleting all nodes in R together with their acceptance conditions and links they are contained in.
2. restricting acceptance conditions of the remaining nodes to the remaining parents.

Definition 11. Let $D = (S, L, C)$ be an ADF. $M \subseteq S$ is *admissible* in D iff there is $R \subseteq S$ such that

1. no element in R attacks an element in M, and
2. M is a stable model of D-R.

M is a *preferred* model of D iff M is (subset) maximal among the sets admissible in D.

Brewka and Woltran also introduced weighted ADFs where an additional weight function w assigns qualitative or numerical weights to the links in the graph. This allows acceptance conditions to be defined in a domain independent way, based on the weights of links rather than on the involved positions. They also showed how the proof standards proposed by Farley and Freeman [7] can be formalized based on this idea.

Since the - rather straightforward - treatment of weights in ADFs will be illustrated in our translation, we do not give further details here and refer the reader to [4].

4. The Translation

We now show how to translate a CAES $\mathcal{S} = \langle \textit{arguments}, \textit{assumptions}, \textit{weights}, \textit{standard} \rangle$ into a dialectical framework $ADF(\mathcal{S})$ (more precisely: a weighted dialectical framework) such that the semantics of $\mathcal{S}$, in other words the outcome of the evaluation, is preserved.

The translation has to take into account that, when scintilla of evidence is used as proof standard, both a proposition p and its complement $\overline{p}$ may be acceptable. For this reason we have to use two nodes for each literal appearing in one of the arguments, one representing the proposition, the other its complement.

We start with the translation of arguments. Let $a = \langle P, E, c \rangle$ be an argument with $P = \{p_1, \ldots, p_k\}$ and $E = \{e_1, \ldots, e_r\}$. The translation of a is the graph (V, R) with

$$V = \{p_1, \ldots, p_k, e_1, \ldots, e_r, c, \overline{c}, a\}$$
$$R = \{(p_i, a) \mid p_i \in P\} \cup \{(e_i, a) \mid e_i \in E\} \cup \{(a, c), (a, \overline{c})\}.$$

Note that it is not sufficient to have nodes representing the literals in the argument. We also need the node a representing the argument itself.[3] We will call the latter type of nodes argument nodes, the other nodes statement nodes.

The translation of all arguments in *arguments* gives us the graph underlying the ADF. We next define the weights associated with the links in the ADF.

Let argument a be as above. The weight function w is defined as follows:

$$\begin{array}{lll} w(x,a) & = + & \text{for } x \in P \\ w(x,a) & = - & \text{for } x \in E \\ w(a,c) & = (+,n) & \text{where } n = \textit{weights}(a) \\ w(a,\overline{c}) & = (-,n) & \text{where } n = \textit{weights}(a) \end{array}$$

Thus, the weights of arguments in Carneades are attached to the links connecting the corresponding argument node with the conclusion and its negation. In addition, pro argument links are marked with a +, con argument links with a −.

[3] One of the reasons for this is that otherwise the resulting ADFs might not be bipolar.

Example 1. Consider the argument

$$a = \langle\{bird\}, \{peng, ostr\}, flies\rangle$$

and assume *weights*$(a) = 0.8$. The ADF graph generated by this argument is shown in Fig. 1 (we mark links with their weights).

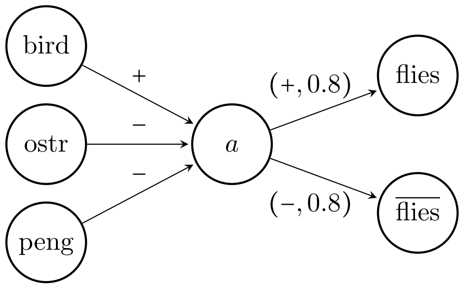

Figure 1. ADF representation of a Carneades argument

The effects of assumptions are represented via the acceptance conditions of argument nodes as follows. Let n be an argument node in the graph obtained from translating a set of arguments as described above. The acceptance condition C_n is defined as:

$C_n(R) = in$ iff
(1) for all p_i with $w(p_i, a) = +$, $p_i \in$ *assumptions* or $\overline{p}_i \notin$ *assumptions* and $p_i \in R$, and
(2) for all e_i with $w(e_i, a) = -$, $p_i \notin$ *assumptions* and $p_i \notin R$ or $\overline{p}_i \in$ *assumptions*.

The acceptance conditions of statement nodes are directly derived from the proof standards for the propositions as specified by the function *standard*. Let m be a statement node, s the proof standard associated to the corresponding proposition via *standard*. We have to distinguish the 5 different cases (we let $max\emptyset = 0$; to avoid repetitions, we attach numbers in brackets to the conditions and use those instead of the conditions):

$s = se$: $C_m(R) = in$ iff [1] for some $r \in R$, $w(r, m) = (+, n)$.
$s = pe$: $C_m(R) = in$ iff [1] and
[2] $max\{n \mid t \in R, w(t, m) = (+, n)\} > max\{n \mid t \in R, w(t, m) = (-, n)\}$.
$s = ce$: $C_m(R) = in$ iff [1] and [2] and
[3] $max\{n \mid t \in R, w(t, m) = (+, n)\} > \alpha$ and
[4] $max\{n \mid t \in R, w(t, m) = (+, n)\} - max\{n \mid t \in R, w(t, m) = (-, n)\} > \beta$.
$s = bd$: $C_m(R) = in$ iff [1] and [2] and [3] and [4] and
[5] $max\{n \mid t \in R, w(t, m) = (-, n)\} < \gamma$
$s = dv$: $C_m(R) = in$ iff [1] and [6] for no $t \in R$, $w(t, m) = (-, n)$.

This concludes our translation. One may observe that, since cycles are not allowed in Carneades, the resulting ADF is acyclic. For such ADFs the semantics we presented earlier coincide:

Proposition 1. *Let the ADF $D = (S, L, C)$ be acyclic. Then D has a single preferred model which coincides with the single stable model and with the well-founded model.*

Proof (sketch): Since D is acyclic, we can show by induction on the number of elements in S that the well-founded model of D, $WF(D)$, is complete in the sense that, for each $s \in S$, $s \notin WF(D)$ implies $s \in E'$, where $(WF(D), E')$ is the least fixpoint of Γ_D. $WF(D)$ is thus the single (two-valued) model of D. Since D has no cycles, there can be no self-supporting links and the least model of the reduct $D^{WF(D)}$ coincides with $WF(D)$. $WF(D)$ is thus a stable model. Since stable models are preferred models, the set of preferred models cannot be empty. We can show, by induction on the number of iterations of the fixpoint operator Γ_D, that the well-founded model is a subset of each preferred model. A similar proof shows that, if $(WF(D), E')$ is the least fixpoint of Γ_D, then no element in E' can be contained in any preferred model of D. Thus, $WF(D)$ is also the single preferred model of D. □

The following result shows that the translation actually preserves the meaning of a CAES.

Proposition 2. *Let $\mathcal{S} = \langle arguments, assumptions, weights, standard\rangle$ be a CAES, $ADF(\mathcal{S}) = (S, L, C)$ the dialectical framework resulting from translating $\mathcal{S}$ as defined above. The following holds:*

1. *An argument $a \in arguments$ is applicable in $\mathcal{S}$ iff the corresponding argument node $a \in S$ is contained in the well-founded (and thus the single preferred and the single stable) model of $ADF(\mathcal{S})$.*
2. *A proposition p is acceptable in $\mathcal{S}$ iff the corresponding statement node $p \in S$ is is contained in the well-founded (and thus the single preferred and the single stable) model of $ADF(\mathcal{S})$.*

Proof (sketch): We prove the proposition by induction on the number n of arguments in $\mathcal{S}$. For $n = 0$ the result is obvious: there is neither an applicable argument nor an acceptable proposition in $\mathcal{S}$, and since $ADF(\mathcal{S})$ is empty, we neither have an argument nor a statement node in the well-founded model of $ADF(\mathcal{S})$, denoted $WF(ADF(\mathcal{S}))$.

Now assume the result holds for n arguments and consider a system with $n + 1$ arguments. Since arguments are acyclic, there must be an argument $a = \langle P, E, c\rangle$ such that c does not appear in the premises or exceptions of any other argument in $\mathcal{S}$. If we disregard this argument we obtain a CAES $\mathcal{S}'$ for which, by induction hypothesis, the proposition holds. For a CAES $\mathcal{S}^*$, let $App(S^*)$ denote the set of applicable arguments in $\mathcal{S}^*$, $Acc(S^*)$ the set of acceptable nodes. Since neither applicability of arguments nor acceptability of propositions in $\mathcal{S}'$ depends on c there are only two cases: either $App(S) = App(S') \cup \{a\}$, or $App(S) = App(S')$. Furthermore, since the acceptability conditions of nodes in $ADF(\mathcal{S}')$ do not depend on c, $WF(ADF(\mathcal{S}'))$ is contained in $WF(ADF(\mathcal{S}'))$. Now, since premises and exceptions of a are in $Acc(S')$ iff they are in the well-founded model of $WF(ADF(\mathcal{S}'))$, and since the acceptance condition for a exactly mirrors the definition of applicability, we have $a \in WF(ADF(\mathcal{S}))$ iff $a \in App(S)$. A similar argument shows that $c \in Acc(S)$ iff $c \in WF(ADF(\mathcal{S}))$. □

So far we have shown that a reconstruction of an acyclic CAES $\mathcal{S}$ as an ADF $ADF(\mathcal{S})$ is indeed possible. Our results also explain why the different Dung semantics do not show up in Carneades: the differences simply do not matter. However, the real advantage of our translation is that we can now lift the restriction of acyclicity, and then, of course, the different semantics do matter. Nothing in our translation hinges on the fact that the set of Carneades arguments is acyclic. Indeed, cycles in the set of arguments of

$\mathcal{S}$ will lead to cycles in $ADF(\mathcal{S})$, yet these cycles are handled - in different ways - by the available semantics of ADFs.

By a *generalized argument evaluation stucture* (GAES) we mean a CAES without the acyclicity restriction for the set of arguments. We define the semantics of GAES as follows:

Definition 12. Let $\mathcal{S} = \langle \textit{arguments}, \textit{assumptions}, \textit{weights}, \textit{standard} \rangle$ be a GAES, $ADF(\mathcal{S}) = (S, L, C)$ the dialectical framework resulting from translating $\mathcal{S}$ as defined above.

1. An argument $a \in$ *arguments* is applicable in $\mathcal{S}$ under grounded (credulous preferred, skeptical preferred, credulous stable, skeptical stable) semantics iff $a \in S$ is contained in the well-founded (all preferred, some preferred, all stable, some stable) model(s) of $ADF(\mathcal{S})$.
2. A proposition p is acceptable in $\mathcal{S}$ under grounded (credulous preferred, skeptical preferred, credulous stable, skeptical stable) semantics iff $p \in S$ is contained in the well-founded (all preferred, some preferred, all stable, some stable) model(s) of $ADF(\mathcal{S})$.

Example 2. Here is a simple example involving a cycle. Assume you are planning your vacation. You plan to go to Greece or to Italy, but cannot afford to visit both countries. Your arguments may thus be:

$$a_1 = \langle \emptyset, \{It\}, Gr \rangle, a_2 = \langle \emptyset, \{Gr\}, It \rangle.$$

These arguments obviously contain a cycle and thus cannot be handled by Carneades. Let n_1 be the weight of a_1, n_2 that of a_2. The translation yields the ADF shown in Fig. 2.

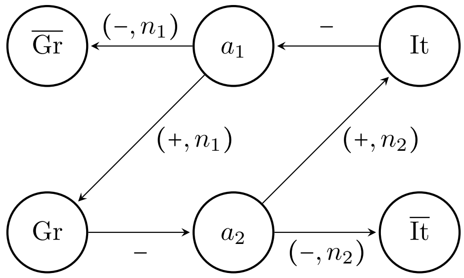

Figure 2. Greece vs. Italy

If we assume that both n_1 and n_2 are greater than α and β, and that further $\gamma > 0$, then the outcome is actually independent of the proof standard chosen for the statement nodes. We get 2 stable models, namely $M_1 = \{a_1, Gr\}$ and $M_2 = \{a_2, It\}$. These models are also the only preferred models. Thus both a_1 and a_2 are applicable under credulous stable and preferred semantics, but neither under skeptical stable nor under skeptical preferred semantics. Similarly, both Gr and It are acceptable under credulous stable and preferred semantics, but neither under skeptical stable nor under skeptical preferred semantics. The well-founded model is empty.

The approach presented here differs significantly from attempts to model proof standards using different Dung semantics, such as [1]. Although the various Dung semantics exhibit different degrees of cautiousness, we have doubts about using these different semantics as a basis for modeling proof standards. First of all, unless modularized variants of argumentation frameworks are used, as in [6] or [3], a chosen semantics is global and doesn't allow different proof standards to be applied to different issues within a single argumentation framework. More importantly, we doubt the various Dung semantics capture the intuitive meanings of legal proof standards. (For a detailed discussion see [9]). Our approach uses proof standards and Dung semantics for different purposes: proof standards are used to aggregate and accrue pro and con arguments about an issue; Dung semantics provide different ways to resolve cyclic arguments.

5. Conclusions

In this paper we have shown that Carneades argument evaluation structures can be reconstructed as abstract dialectical frameworks. This has several benefits, both from the point of view of ADFs and from the point of view of Carneades.

1. It shows that ADFs not only generalize Dung argumentation frameworks - which have been the starting point for their development. They also generalize Carneades argument evaluation structures. This provides evidence that the ADF framework is indeed a useful tool in the theory of argumentation.
2. It clarifies the relationship between Carneades and Dung AFs, showing that both are instances of ADFs. It thus helps to put Carneades on an equally solid formal foundation.
3. Finally, as we have seen, it allows us to lift the restriction of Carneades to acyclic argument structures.

References

[1] Katie Atkinson and Trevor J. M. Bench-Capon. Argumentation and standards of proof. In *Proc. ICAIL*, pages 107–116, 2007.

[2] Stefan Ballnat and Thomas Gordon. Goal selection in argumentation processes. In *Proc. COMMA* (this volume), 2010.

[3] Gerhard Brewka and Thomas Eiter. Argumentation context systems: A framework for abstract group argumentation. In *LPNMR*, pages 44–57, 2009.

[4] Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Proc. Principles of Knowledge Representation and Reasoning*, pages 102–111, 2010.

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

[6] Phan Minh Dung, Phan Minh Thang, and Nguyen Duy Hung. Modular argumentation for modelling legal doctrines of performance relief. In *ICAIL '09: Proceedings of the 12th International Conference on Artificial Intelligence and Law*, pages 128–136, New York, NY, USA, 2009. ACM.

[7] Arthur M. Farley and Kathleen Freeman. Burden of proof in legal argumentation. In *Proc. ICAIL'95*, pages 156–164, 1995.

[8] Thomas F. Gordon, Henry Prakken, and Douglas Walton. The Carneades model of argument and burden of proof. *Artif. Intell.*, 171(10-15):875–896, 2007.

[9] Thomas F. Gordon and Douglas Walton. Proof burdens and standards. In Iyad Rahwan and Guillermo Simari, editors, *Argumentation in Artificial Intelligence*, pages 239–258. 2009.

[10] Matthias Grabmair, Thomas Gordon, and Douglas Walton. Probabilistic semantics for the Carneades argument model using Bayesian belief networks. In *Proc. COMMA* (this volume), 2010.