Computational Models of Argument
H. Prakken et al. (Eds.)

© 2020 The authors and IOS Press.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/FAIA200493

Computing Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks with Recursive Attack and Support Relations

Gianvincenzo ALFANO [1], Sergio GRECO and Francesco PARISI
DIMES Department, University of Calabria, Italy

Abstract. Attack-Support Argumentation Framework (ASAF) is an extension of the Bipolar Argumentation Framework that allows for attacks and supports not only between arguments but also targeting attacks and supports at any level. In this paper we propose an incremental approach for computing the skeptical preferred acceptance in dynamic ASAFs. Specifically, we investigate how the skeptical acceptance of a goal element (an argument, an attack, or a support) evolves when a given ASAF is updated by adding or retracting an argument, an attack, or a support, and propose an incremental algorithm for solving this problem. Our approach relies on identifying a portion of the given ASAF which is sufficient to determine the status of the goal w.r.t. the updated ASAF. We experimentally evaluate our approach showing that it outperforms the computation from scratch on average.

Keywords. Abstract argumentation, Higher-order interactions, Incremental computation

1. Introduction

Formal argumentation has emerged as one of the important fields in Artificial Intelligence [17,37,11]. In particular, Dung's framework is a simple, yet powerful formalism for modelling disputes between two or more agents [29]. An abstract Argumentation Framework (AF) consists of a set of *arguments* and a binary *attack* relation over the set of arguments that specifies the *interactions* between arguments: intuitively, if argument a attacks argument b, then b is acceptable only if a is not. Hence, arguments are abstract entities whose role is entirely determined by the interactions specified by the attack relation.

Dung's framework has been extended in many different ways, including the introduction of new kinds of interactions between arguments and/or attacks. In particular, the Bipolar Argumentation Framework is an interesting extension of the Dung's framework which allows for modelling the *support* between arguments [9,27,39]. Further extensions consider second-order interactions [20], e.g., attacks to attacks/supports, as well as more general forms of interactions such as Argumentation Frameworks with Recursive Attacks

[1]Corresponding Author: Gianvincenzo Alfano, e-mail: g.alfano dimes.unical.it

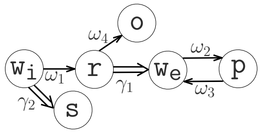

Figure 1. ASAF of Example 1

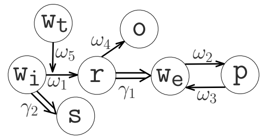

Figure 2. ASAF for winter scenario

[13] and Attack-Support Argumentation Frameworks (ASAFs) [31], where attacks and supports can be recursively attacked or supported.

Example 1. *Consider a scenario to decide whether to play tennis. Assume we have the following arguments:* $\mathtt{w_i}$ *(it is windy),* $\mathtt{r}$ *(it rains),* $\mathtt{w_e}$ *(the court is wet),* $\mathtt{p}$ *(play tennis),* $\mathtt{s}$ *(need a sweatshirt),* $\mathtt{o}$ *(tennis racket shop is open), and the implications: (ω_1) if it is windy, then it does not rain, (ω_2) if the court is wet, then we cannot play tennis, (ω_3) if we play tennis then the court is not wet, (ω_4) if it rains then the tennis racket shop is not open, (γ_1) if it rains, then the court is wet, and (γ_2) if it is windy, then we need a sweatshirt. This situation can be modeled by using the ASAF of Figure 1, where ω_1, ω_2 and ω_3 are* attacks *(denoted by $\rightarrow$), and γ_1 and γ_2 are* supports *(denoted by $\Rightarrow$).* □

Several interpretations of the notion of support have been proposed [25,27]. The *necessary* support [13,36] adopted in ASAF is intended to capture the following intuition: if a supports b, then the acceptance of a is necessary to get the acceptance of b; equivalently, accepting b implies accepting a. The meaning of an ASAF is given by extensions which also include attacks and supports that contribute to determine the set of accepted arguments. For instance, considering the well-known *preferred* semantics—one of the most popular argumentation semantics [22]—the framework of Figure 1 has a unique extension, that is the set $\{\mathtt{w_i}, \mathtt{s}, \mathtt{p}, \mathtt{o}, \omega_1, \omega_3, \gamma_1, \gamma_2\}$.

However, in practice, argumentation frameworks can be dynamic systems [12,15,16,18,26,34]. In fact, typically an ASAF represents a temporary situation, and new arguments, attacks and supports (at any level) can be added/removed to take into account new available knowledge. For instance, in our running example, assume now that there exists also an argument $\mathtt{w_t}$ (we are in the winter season) that attacks ω_1 (in the winter season ω_1 cannot be applied). The updated scenario can be modeled by an ASAF shown in Figure 2 where the new attack is labelled as ω_5 (ω_5 is an example of second-level attack).

Recently, there has been a growing interest in studying dynamics of different argumentation systems, considering the Dung framework [2,5,15,19,28], Bipolar AF and AF with second order attacks [3,4], ASAF [1], and structured argumentation formalisms [7,8]. This is motivated by the fact that most of the argumentation problems have high computational complexity [30,33]. In particular, skeptical reasoning under the preferred semantics is in the second level of the polynomial hierarchy. However, in practice, incremental computation techniques could improve performance, as they only require to reconsider the acceptance status of those arguments and interactions that are affected by the new information.

In this paper we propose an incremental approach for computing the skeptical preferred acceptance of a goal element of an ASAF after performing an update. Specifically, we propose a technique addressing the following problem: given an ASAF Δ, a goal element G whose skeptical preferred acceptance w.r.t. Δ is known, and an update u consisting of the addition/removal of an argument/attack/support, decide whether G is skeptically preferred accepted w.r.t. the updated ASAF $u(\Delta)$, that is, decide if G belongs to every preferred extension of $u(\Delta)$.

Contributions. We make the following contributions:

- Given an update and a goal element (an argument, an attack, or a support), we identify a set of elements, called *alterable set*, which contains the elements whose acceptance status may change after the update and propagate up to the goal.
- Given the alterable set, we define the *Proxy ASAF* that allows us to compute the skeptical preferred acceptance of a goal by focusing on a (potentially smaller) ASAF containing the alterable set as well as additional elements and interactions needed to determine the status of the elements in the alterable set.
- We introduce an incremental algorithm for computing the skeptical preferred acceptance of a goal within a dynamic ASAF. It enables the computation on the Proxy ASAF, provided that an external solver for ASAFs is given.
- Since to the best of our knowledge there is no available solver for the direct computation on ASAF, we propose a version of the algorithm that, using a translation of our problem to the AF domain, allows us to use any (non-incremental) state-of-the-art AF solver to compute the skeptical preferred acceptance for ASAFs
- We provide an experimental analysis showing the effectiveness of our approach.

To the best of our knowledge, this is the first paper addressing the problem of efficiently and incrementally computing skeptical acceptance for dynamic ASAFs.

2. Preliminaries

We start by briefly reviewing the Dung's framework [29] and the Attack-Support Argumentation Framework (ASAF) (for a full presentation of ASAF see [31]).

An abstract *Argumentation Framework* (AF) is a pair $\langle \mathbb{A}, \Sigma \rangle$, where $\mathbb{A}$ is a set of *arguments* and $\Sigma \subseteq \mathbb{A} \times \mathbb{A}$ is a set of *attacks*. An AF can be seen as a directed graph, whose nodes represent arguments and edges represent attacks.

Given an AF $\Lambda = \langle \mathbb{A}, \Sigma \rangle$ and a set $S \subseteq \mathbb{A}$ of arguments, an argument $a \in \mathbb{A}$ is said to be *i*) *attacked* (or, equivalently, *defeated*) w.r.t. S iff $\exists b \in S$ such that $(b, a) \in \Sigma$, and *ii*) *acceptable* w.r.t. S iff for every argument $b \in \mathbb{A}$ with $(b, a) \in \Sigma$, there is $c \in S$ such that $(c, b) \in \Sigma$. The sets of defeated and acceptable arguments w.r.t. S can be defined as follows (where Λ is understood):

- $Def(S) = \{a \in \mathbb{A} \mid \exists\, b \in S \,.\, (b, a) \in \Sigma\}$;
- $Acc(S) = \{a \in \mathbb{A} \mid \forall\, b \in \mathbb{A} \,.\, (b, a) \in \Sigma \;\Rightarrow\; b \in Def(S)\}$.

Given an AF $\langle \mathbb{A}, \Sigma \rangle$, a set $S \subseteq \mathbb{A}$ of arguments is said to be: (*i*) *conflict-free* iff $S \cap Def(S) = \emptyset$; (*ii*) *admissible* iff it is conflict-free and $S \subseteq Acc(S)$. Moreover, $S \subseteq \mathbb{A}$ is said to be a a *preferred extension* iff it is conflict-free, $S = Acc(S)$, and maximal (w.r.t. $\subseteq$). The set of preferred extensions of an AF Λ will be denoted by $\mathcal{PR}(\Lambda)$.

Example 2. *Let* $\Lambda = \langle \mathbb{A}, \Sigma \rangle$ *be an AF where* $\mathbb{A} = \{\mathtt{r}, \mathtt{w_e}, \mathtt{p}\}$ *and* $\Sigma = \{(\mathtt{w_e}, \mathtt{p}), (\mathtt{p}, \mathtt{w_e})\}$. *The set of preferred extensions is* $\mathcal{PR}(\Lambda) = \{\{\mathtt{r}, \mathtt{w_e}\}, \{\mathtt{r}, \mathtt{p}\}\}$. □

Given an AF $\Lambda = \langle \mathbb{A}, \Sigma \rangle$ and an argument $G \in \mathbb{A}$, we say that G is skeptically preferred accepted w.r.t. Λ iff for each preferred extension E of Λ it holds that $G \in E$. For instance, for the AF in Example 2, we have that $\mathtt{r}$ is skeptically preferred accepted.

Attack-Support Argumentation Framework

An *Attack-Support Argumentation Framework (ASAF)* [31] ASAF is a triple $\langle A, \Omega, \Gamma\rangle$, where A is a set of arguments, $\Omega \subseteq A \times (A \cup \Omega \cup \Gamma)$ is a set of attacks, and $\Gamma \subseteq A \times (A \cup \Omega \cup \Gamma)$ is a set of supports. It is assumed that Γ is acyclic and $\Omega \cap \Gamma = \emptyset$.

In the following, given an ASAF $\langle A, \Omega, \Gamma\rangle$, to simplify the notation, we use symbols ω_i (or simply ω) to denote attacks (e.g., $\omega = (a, X) \in \Omega$) and symbols γ_i (or simply γ) to denote supports (e.g., $\gamma = (b, Y) \in \Gamma$); we also use δ to denote an element in $\Omega \cup \Gamma$. Moreover, given an ASAF $\langle A, \Omega, \Gamma\rangle$, for any attack or support $\delta = (a, Y) \in \Omega \cup \Gamma$, we use $\mathtt{s}(\delta)$ and $\mathtt{t}(\delta)$ to denote, respectively, the source argument a and the target element Y of δ. Note that Y can be an argument, an attack, or a support. Attacks and supports whose target is an argument are said to be *first-level* interactions, while attacks and supports whose target is an interaction of level i are said to be interactions of level $i + 1$.

An ASAF Δ can be represented by a graph-like structure $\mathcal{G}_\Delta$ where an argument $a \in A$ is a node in $\mathcal{G}_\Delta$, an attack $\omega = (a, X) \in \Omega$ is graphically denoted as an edge $a \xrightarrow{\omega} X$ in $\mathcal{G}_\Delta$, and a support $\gamma = (b, Y) \in \Gamma$ is graphically denoted as an edge $b \stackrel{\gamma}{\Longrightarrow} Y$ in $\mathcal{G}_\Delta$. For instance, the graph in Figure 1 represents the ASAF of Example 1, that is, an ASAF $\Delta = \langle A, \Omega, \Gamma\rangle$, where $A = \{\mathtt{w_i}, \mathtt{r}, \mathtt{s}, \mathtt{o}, \mathtt{w_e}, \mathtt{p}\}$, $\Omega = \{\omega_1 = (\mathtt{w_i}, \mathtt{r}), \omega_2 = (\mathtt{w_e}, \mathtt{p}), \omega_3 = (\mathtt{p}, \mathtt{w_e}), \omega_4 = (\mathtt{r}, \mathtt{o})\}$, and $\Gamma = \{\gamma_1 = (\mathtt{r}, \mathtt{w_e}), \omega_2 = (\mathtt{w_i}, \mathtt{s})\}$.

Attacks and supports in an ASAF can also be attacked and supported, and extensions may contain arguments, attacks and supports. The semantics proposed in [31] combines the interpretation of attacks of Argumentation Frameworks with Recursive Attacks [13] with that of Bipolar AFs with necessary support [25], as formalized in what follows.

Given an ASAF $\langle A, \Omega, \Gamma\rangle$, a *support path* $a_0 \stackrel{+}{\Rightarrow} X$ from a_0 to X is is a sequence of n supports $a_0 \stackrel{\gamma_1}{\Longrightarrow} a_1 \stackrel{\gamma_2}{\Longrightarrow} \ldots a_{n-1} \stackrel{\gamma_n}{\Longrightarrow} X$, where each a_i (with $0 \leq i < n$) is an argument and X is either an argument, an attack, or a support. We use $\Gamma^+ = \{(a, X) \mid a \in A,\ X \in (A \cup \Omega \cup \Gamma),\ a \stackrel{+}{\Rightarrow} X\}$ to denote the set of pairs (a, X) such that there exists a (not empty) support path from a to X.

Given an element $X \in (A \cup \Omega \cup \Gamma)$ and an attack $\omega \in \Omega$, we say that ω *(directly or indirectly) attacks* X (denoted by ω $\mathtt{def}$ X) if either $\mathtt{t}(\omega) = X$ or $\mathtt{t}(\omega) = \mathtt{s}(X)$. Moreover, given a set $S \subseteq A \cup \Omega \cup \Gamma$, we say that ω *extendedly defeats* X *given* S (denoted as ω $\mathtt{def_S}$ X) if either ω $\mathtt{def}$ X or there exists $b \in A$ such that $\mathtt{t}(\omega) = b$ and either $(b, X) \in (\Gamma \cap S)^+$ or $(b, s(X)) \in (\Gamma \cap S)^+$. For any ASAF Δ and $S \subseteq A \cup \Omega \cup \Gamma$, the *defeated* and *acceptable* sets (given S) are:

- $Def(S) = \{X \in A \cup \Omega \cup \Gamma \mid \exists\, \omega \in \Omega \cap S \,.\, \omega\ \mathtt{def_S}\ X\}$
- $Acc(S) = \{X \in A \cup \Omega \cup \Gamma \mid \forall \omega \in \Omega \,.\, \omega\ \mathtt{def_S} X \Rightarrow \omega \in Def(S)\}$.

The notions of *conflict-free*, *admissible sets*, and the *preferred extensions* for ASAF can be defined as done earlier (before Example 2) for the AF but considering $S \subseteq A \cup \Omega \cup \Gamma$ and by using the definitions of defeated and acceptable sets reported above.

Finally, given an ASAF $\Delta = \langle A, \Omega, \Gamma\rangle$ and an element $G \in A \cup \Omega \cup \Gamma$, we say that G is skeptically preferred accepted w.r.t. Δ iff for each preferred extension E of Δ it holds that $G \in E$. In the following, we use $SA_\Delta(G)$ to denote the skeptical preferred acceptance (either *true* or *false*) of G w.r.t. ASAF Δ.

Example 3. *Let* $\Delta = \langle\{\mathtt{w_i}, \mathtt{r}, \mathtt{s}, \mathtt{o}, \mathtt{w_e}, \mathtt{p}, \mathtt{w_t}\}, \{\omega_1 = (\mathtt{w_i}, \mathtt{r}), \omega_2 = (\mathtt{w_e}, \mathtt{p}), \omega_3 = (\mathtt{p}, \mathtt{w_e}), \omega_4 = (\mathtt{r}, \mathtt{o}), \omega_5 = (\mathtt{w_t}, (\mathtt{w_i}, \mathtt{r}))\}, \{\gamma_1 = (\mathtt{r}, \mathtt{w_e}), \gamma_2 = (\mathtt{w_i}, \mathtt{s})\}\rangle$ *be the ASAF of Figure 2. The set of preferred extensions of* Δ *is* $\mathcal{PR}(\Delta) = \{\{\mathtt{w_i}, \mathtt{r}, \gamma_1, \mathtt{s}, \mathtt{w_e}, \omega_2, \mathtt{w_t}, \omega_4, \omega_5, \gamma_2\},$

$\{\mathtt{w_i}, \mathtt{r}, \gamma_1, \mathtt{s}, \mathtt{p}, \omega_3, \mathtt{w_t}, \omega_4, \omega_5, \gamma_2\}\}$. *Thus, the set of elements X of Δ that are skeptically accepted (i.e., those for which $SA_\Delta(X) = true$) is* $\{\mathtt{w_i}, \mathtt{r}, \gamma_1, \mathtt{s}, \mathtt{w_t}, \omega_4, \omega_5, \gamma_2\}$. □

Updates for ASAF

An *update* consists of the addition (resp., removal) of an attack or a support not present (resp., present) in a given ASAF, as next formalized.

Definition 1 (Update for ASAF)**.** *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ an ASAF, and $\delta \in A \times (A \cup \Omega \cup \Gamma)$. An* update *$u$ over Δ is of one of the forms below, and when applied to Δ yields the updated ASAF $u(\Delta) = \langle A, \Omega', \Gamma'\rangle$, with Ω' and Γ' defined as follows:*

- *$u = +\delta$ where $\delta \notin (\Omega \cup \Gamma)$. If δ is an attack, then $\Omega' = \Omega \cup \{\delta\}$ and $\Gamma' = \Gamma$, otherwise $\Omega' = \Omega$, $\Gamma' = \Gamma \cup \{\delta\}$ and Γ' is acyclic.*
- *$u = -\delta$ where $\delta \in \Omega \cup \Gamma$ and there is no $\delta' \in \Omega \cup \Gamma$ such that $\mathbf{t}(\delta') = \delta$. In this case, $\Omega' = \Omega \setminus \{\delta\}$ and $\Gamma' = \Gamma \setminus \{\delta\}$.*

In the following, for simplicity, we write $\pm\delta$ for the addition or removal of an attack or a support $(\mathbf{s}(\delta), \mathbf{t}(\delta))$. Then, for an update $u = +\delta$, the interaction δ must not belong to the attack and support relations of the ASAF it will be applied on, and the source and target of δ must belong to the ASAF; moreover, the support relation of the updated ASAF must remain acyclic. Moreover, for an update $u = -\delta$, the interaction δ cannot be targeted by any other interaction in the ASAF.

As for an update u consisting of the addition (resp. deletion) of a set of *isolated* arguments (i.e., arguments not connected to any other element in the graph), it is easy to see that if $u(\Delta)$ is obtained from Δ through the addition (resp. deletion) of a set S of isolated argument, then every argument in S is trivially skeptically preferred accepted (resp., not accepted) w.r.t. $u(\Delta)$. Indeed, if E is an extension for Δ, then $E' = E \cup S$ (resp. $E' = E \setminus S$) is an extension for $u(\Delta)$ containing every (resp., none) argument in S. Of course, if arguments in S are not isolated, for addition we can first add isolated arguments and then add interactions (attacks or supports) involving these arguments, while for deletion we can first delete all interactions involving arguments in S and then delete isolated arguments. Thus we do not consider these kinds of updates in the following, and w.l.o.g. focus on updates consisting of the addition or deletion of an attack or a support.

3. Incremental Computation of Skeptical Preferred Acceptance

In this section, given an ASAF and an update for it, we propose an incremental technique for computing the skeptical preferred acceptance of a given goal element.

First we identify a set of *alterable* elements, that is, a set of arguments, attacks, and supports whose acceptance status may change after performing an update, and such that the change may impact on the acceptance status of the goal. We start by defining the set of elements that are reachable from a given element X of an ASAF. This set includes X and its *neighbors*, i.e. the target of X and, if X is an argument, also the interactions originating from X and the targets of such interactions, as formalized in what follows.

Definition 2 (Set of neighbors)**.** *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF. The set $N_\Delta(X)$ of neighbors of an element $X \in A \cup \Omega \cup \Gamma$ is:*
i) $\{X, \mathbf{t}(X)\}$ if $X \in \Omega \cup \Gamma$,   *ii) $\{X\} \cup \{Y, \mathbf{t}(Y) \mid X = \mathbf{s}(Y),\ Y \in \Omega \cup \Gamma\}$ if $X \in A$.*

For instance, for the ASAF Δ of Figure 2, we have that $N_\Delta(\mathtt{w_i}) = \{\mathtt{w_i}, \omega_1, \mathtt{r}, \gamma_2, \mathtt{s}\}$, $N_\Delta(\omega_5) = \{\omega_5, \omega_1\}$, and $N_\Delta(\omega_1) = \{\omega_1, \mathtt{r}\}$. The set of elements that are reachable from X consists of $N_\Delta(X)$ plus the elements which are reachable from $N_\Delta(X)$.

Definition 3 (Reachable elements)**.** *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF. Given $X, Y \in A \cup \Omega \cup \Gamma$ we say that Y is reachable from X in Δ iff either i) $Y \in N_\Delta(X)$ or ii) $\exists\, Z \in A \cup \Omega \cup \Gamma$ such that $Z \in N_\Delta(X)$ and Y is reachable from Z in Δ.*
We use $Reach_\Delta(X)$ to denote the set of elements of Δ that are reachable from X in Δ.

For the ASAF Δ of Figure 2, $Reach_\Delta(\omega_5) = \{\omega_5, \omega_1, \mathtt{r}, \omega_4, \gamma_1, \mathtt{o}, \mathtt{w_e}, \omega_2, \mathtt{p}, \omega_3\}$.

In the following, we use Δ^u to denote the larger ASAF between Δ and $u(\Delta)$, that is, Δ^u is *i*) the updated ASAF $u(\Delta)$ if u is an addition update (it includes the interaction added through u), *ii*) the original ASAF Δ if u is a deletion update (the removed interaction is also considered in Δ^u).

We are now ready to define the alterable set for an ASAF w.r.t. a given update.

Definition 4 (Alterable Set)**.** *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF, $u = \pm\delta$ an update, and $G \in A \cup \Omega \cup \Gamma$ a (goal) element. Let*

- $Alt_0(\Delta, u, G) = \begin{cases} \emptyset & if\ G \notin Reach_{\Delta^u}(\delta); \\ N_{\Delta^u}(\delta) & otherwise. \end{cases}$
- $Alt_{i+1}(\Delta, u, G) = Alt_i(\Delta, u, G) \cup \{Z \mid Z \in N_{\Delta^u}(Y),\ Y \in Alt_i(\Delta, u, G),\ G \in Reach_{\Delta^u}(Z)\}.$

Let n be the natural number such that $Alt_n(\Delta, u, G) = Alt_{n+1}(\Delta, u, G)$. Then alterable set $Alt(\Delta, u, G)$ is $Alt_n(\Delta, u, G)$.

Thus, the alterable set is iteratively defined by $n+1$ steps (with $n \le |A|+|\Omega|+|\Gamma|$), each of them consisting of the addition of at least a neighbor of an element in the set built at the previous step and allowing to reach the goal G. It is easy to see that, for any element G, it is the case that $Alt(\Delta, u, G) \subseteq Reach_{\Delta^u}(\delta)$, where $u = \pm\delta$.

Example 4. *Consider the ASAF Δ of Figure 2, the update $u = -\omega_5$, and assume that $\mathtt{p}$ is the goal element. Note that, differently from the introduction, the update considered here is a deletion. Then, $Alt_0(\Delta, u, \mathtt{p}) = \{\omega_5, \omega_1\}$ as $\mathtt{p} \in Reach_{\Delta^u}(\omega_5)$. $Alt_1(\Delta, u, \mathtt{p}) = Alt_0(\Delta, u, \mathtt{p}) \cup \{\mathtt{r}\}$, $Alt_2(\Delta, u, \mathtt{p}) = Alt_1(\Delta, u, \mathtt{p}) \cup \{\gamma_1, \mathtt{w_e}\}$ (herein, ω_4 and $\mathtt{o}$ are not included as they do not allow to reach the goal in Δ^u), $Alt_3(\Delta, u, \mathtt{p}) = Alt_2(\Delta, u, \mathtt{p}) \cup \{\omega_2, \mathtt{p}\}$. Finally, $Alt_4(\Delta, u, \mathtt{p}) = Alt_3(\Delta, u, \mathtt{p}) \cup \{\omega_3\}$, and thus $Alt(\Delta, u, \mathtt{p}) = \{\omega_5, \omega_1, \mathtt{r}, \gamma_1, \mathtt{w_e}, \omega_2, \mathtt{p}, \omega_3\} \subseteq Reach_{\Delta^u}(\omega_5)$.*

The following theorem states that, after performing an update, the skeptical preferred acceptance of an element does not change if the alterable set is empty.

Theorem 1. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF, u an update, $u(\Delta)$ the updated ASAF, and G a goal element in $A \cup \Omega \cup \Gamma$. Therefore, if $Alt(\Delta, u, G) = \emptyset$ then $SA_{u(\Delta)}(G) = SA_\Delta(G)$.*

If the alterable set is not empty, we identify a (potentially small) portion of the given ASAF, called *Proxy ASAF*, that is sufficient to perform the computation of the skeptical preferred acceptance of the goal without considering the entire ASAF.

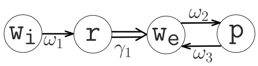

Figure 3. Proxy ASAF

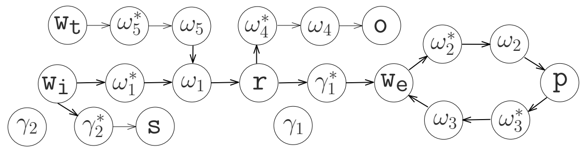

Figure 4. AF for the ASAF of Figure 2

Before defining the Proxy ASAF, we introduce some notation. Given an ASAF $\Delta = \langle A, \Omega, \Gamma\rangle$, for a set $S \subseteq A \cup \Omega \cup \Gamma$ of elements of Δ, we use $Reach^{-1}_{\Delta}(S) = \{Y \in A \cup \Omega \cup \Gamma \mid \ X \in S, X \in Reach_{\Delta}(Y)\}$ to denote the set of elements from which the elements in S are reachable in Δ. Moreover, we use $\Delta{\downarrow_S}$ to denote the *restriction* of an ASAF $\Delta = \langle A, \Omega, \Gamma\rangle$ to a set S of elements, that is $\Delta{\downarrow_S} = \langle A_S, \Omega_S, \Gamma_S\rangle$, where $A_S = A \cap S$, $\Omega_S = \{\omega \in \Omega \mid \mathbf{s}(\omega) \in A_S \wedge \mathbf{t}(\omega) \in (A_S \cup \Omega_S \cup \Gamma_S)\}$, and $\Gamma_S = \{\gamma \in \Gamma \mid \mathbf{s}(\gamma) \in A_S \wedge \mathbf{t}(\gamma) \in (A_S \cup \Omega_S \cup \Gamma_S)\}$.

The Proxy ASAF is the restriction of the updated ASAF $u(\Delta)$ to the alterable set plus the elements of $u(\Delta)$ that can reach an element in that set.

Definition 5 (Proxy ASAF). *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF, $u = \pm\delta$ an update, and $G \in A \cup \Omega \cup \Gamma$ a goal element. Let $S = Alt(\Delta, u, G)$. The Proxy ASAF of Δ w.r.t u and G is $PASAF(\Delta, u, G) = u(\Delta){\downarrow_{S \cup Reach^{-1}_{u(\Delta)}(S)}}$.*

Example 5. *Continuing from Example 4, the Proxy ASAF $PASAF(\Delta, u, \mathtt{p})$ is given by considering the restriction of the updated ASAF $u(\Delta)$ to the alterable set $S = Alt(\Delta, u, \mathtt{p})$ union $Reach^{-1}_{u(\Delta)}(S) = \{\mathtt{w_i}\}$, as reported in Figure 3.*

Observe that $PASAF(\Delta, u, G)$ is empty if $Alt(\Delta, u, G)$ is empty. In this case we can use the result of Theorem 1 to compute the skeptical acceptance. In contrast, the following theorem tells us how to use the Proxy ASAF to compute the skeptical preferred acceptance when the alterable set is not empty.

Theorem 2. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF, u an update, $u(\Delta)$ the updated ASAF, and a goal element $G \in A \cup \Omega \cup \Gamma$. If $Alt(\Delta, u, G) \neq \emptyset$ then G is skeptically preferred accepted w.r.t. $u(\Delta)$ iff it is skeptically preferred accepted w.r.t. the Proxy ASAF $PASAF(\Delta, u, G)$.*

Example 6. *Continuing from Example 5, $\mathtt{p}$ is skeptically preferred accepted w.r.t. the ASAF $u(\Delta)$ since $\mathtt{p}$ is skeptically preferred accepted w.r.t. the Proxy ASAF $PASAF(\Delta, u, \mathtt{p})$ of Figure 3 whose unique preferred extension is $\{\mathtt{w_i}, \mathtt{p}, \omega_1, \gamma_1, \omega_3\}$.*

3.1. Incremental Algorithm

The results of Theorems 1 and 2 allow us to define Algorithm 1 to decide the skeptical preferred acceptance of an element G w.r.t. an ASAF Δ updated by $u = \pm\delta$. Given the initial skeptical preferred acceptance $SA_{\Delta}(G)$, the skeptical preferred acceptance $SA_{u(\Delta)}(G)$ w.r.t. the updated ASAF is incrementally computed, thus enabling consecutive invocations of the algorithm to perform sequences of updates. Algorithm 1 works as follows. First the alterable set is computed at Line 1. Using result of Theorem 1, if the alterable set is empty then the acceptance status of G does not change after the update, and the algorithm returns the initial status at Line 3. Otherwise, the Proxy ASAF

Algorithm 1 *ASAF-SA*($\Delta, u, G, SA_\Delta(G)$, ASAF-Solver)

Input: ASAF $\Delta = \langle A, \Omega, \Gamma \rangle$, update u, goal $G \in A \cup \Omega \cup \Gamma$, initial skeptical preferred acceptance $SA_\Delta(G)$, function ASAF-Solver computing the skeptical preferred acceptance of a goal element for an ASAF.

Output: updated skeptically preferred acceptance of G w.r.t $u(\Delta)$.

1: Let $S = Alt(\Delta, u, G)$
2: **if** $S = \emptyset$ **then**
3:     **return** $SA_\Delta(G)$;
4: Let $\Delta_P = PASAF(\Delta, u, G)$
5: **return** ASAF-Solver(G, Δ_P)

is built at Line 5 and, using Theorem 2, the skeptical acceptance of G can be computed by invoking an external ASAF-Solver that decides whether G is skeptically accepted by performing the computation on the Proxy ASAF (Line 5).

Algorithm 1 assumes that an ASAF solver is given. That is, in principle, our approach enables any external solver for ASAF to be used for the incremental computation of the preferred skeptical acceptance. However, to the best of our knowledge, currently there is no solver that directly performs the computation of skeptical acceptance on ASAFs (this is also due to the fact that the ASAF proposal is a quite recent, compared to Dung's framework for which several solvers have become available during the last few years). Therefore, instead of performing the computation on the Proxy ASAF, we leverage on a transformation of the Proxy ASAF to a Dung's framework to eventually compute the skeptical acceptance of the given goal. This makes our approach working with any available AF solver for the computation of the skeptical preferred acceptance. As explained below, we use the meta-AF approach recently proposed in [1] for computing ASAF extensions that can be adopted also for our scope.

Enabling the computation at the AF level

In this section, we first briefly review the transformation presented in [1] that allow us to characterize an ASAF in terms of an AF whose extensions (under preferred, grounded, complete, and stable semantics) are in a one-to-one correspondence with those of the given ASAF. Then, we show how to use this result to compute the skeptical acceptance.

An *AF for an ASAF* is an AF that encodes every argument, attack, and support of the given ASAF. The set of arguments of the AF consists of the arguments of Δ plus a pair of arguments, ω and ω^*, for each attack ω in Δ and a pair of arguments, γ and γ^*, for each support γ in Δ. Arguments ω and ω^* determine whether ω is accepted or not, and are used to propagate defeats on the source of ω to the attack itself. Argument γ represents the support itself and is used to determine whether it is accepted or not, whereas argument γ^* is used to propagate defeats on the source of γ to its target. Then, the attacks of the AF are as follows. For each attack ω in Δ, the AF contains a chain of 3 attacks starting in the source of ω and ending in its target, with intermediate arguments ω^* and ω; moreover, if the target of ω is a support γ, then an attack between ω and both γ^* and γ is added to the AF. For each support γ in Δ, the AF contains a chain of 2 attacks starting in the source of γ and ending in its target, with intermediate argument γ^*; finally, if the target of γ is a support γ_1, an attack between γ^* and γ_1^* is added.

Definition 6 (AF for ASAF [1]). *Let* $\Delta = \langle A, \Omega, \Gamma \rangle$ *be an ASAF. The* AF for Δ *is* $\Lambda_\Delta = \langle \mathbb{A}_\Delta, \Sigma_\Lambda \rangle$, *where:*

- $\mathbb{A}_\Delta = A \cup \{\omega, \omega^* \mid \omega \in \Omega\} \cup \{\gamma, \gamma^* \mid \gamma \in \Gamma\}$.
- $\Sigma_\Lambda = \{(\mathbf{s}(\omega), \omega^*), (\omega^*, \omega), (\omega, \mathbf{t}(\omega)) \,|\, \omega \in \Omega\} \cup \{(\omega, \mathbf{t}(\omega)^*) \mid \omega \in \Omega,\, \mathbf{t}(\omega) \in \Gamma\}$ $\cup \{(\mathbf{s}(\gamma), \gamma^*), (\gamma^*, \mathbf{t}(\gamma)) \mid \gamma \in \Gamma\} \cup \{(\gamma^*, \mathbf{t}(\gamma)^*) \mid \gamma \in \Gamma, \mathbf{t}(\gamma) \in \Gamma\}$.

Example 7. *The AF for the ASAF Δ of Figure 2 is Λ_Δ shown in Figure 4. For instance, the attack $\omega_1 = (\mathtt{w_i}, \mathtt{r})$ corresponds to the chain of attacks from $\mathtt{w_i}$ to $\mathtt{r}$ through ω_1 and ω_1^*, while $\omega_5 = (\mathtt{w_t}, \omega_1)$ corresponds to the attacks $(\mathtt{w_t}, \omega_5^*), (\omega_5^*, \omega_5), (\omega_5, \omega_1)$.*

In [1], it is shown that there exists a one-to-one correspondence between the preferred extensions of an ASAF Δ and the preferred extensions of the AF Λ_Δ for Δ, modulo meta-arguments ω^* and γ^*. This equivalence between extensions of an ASAF and extensions of the corresponding AF allow us to state the following result.

Theorem 3. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be an ASAF, Λ_Δ the AF for Δ, and G an element in $A \cup \Omega \cup \Gamma$. Therefore, G is skeptically preferred accepted w.r.t. Δ iff the argument corresponding to G is skeptically preferred accepted w.r.t. Λ_Δ.*

Algorithm 2: a variant of Algorithm 1 using an AF solver. To perform the computation of the skeptical preferred acceptance by using a state-of-the-art AF solver, we modify Algorithm 1 as follows. Let ASAFtoAF be a function that takes as input an ASAF Δ and returns the corresponding AF Λ_Δ. Then, the invocation of the ASAF solver at Line 5 of Algorithm 1 is replaced by AF-Solver($\overline{G}$, ASAFtoAF(Δ_P)), where AF-Solver is a function computing the skeptical preferred acceptance of a given argument w.r.t. a given AF, and $\overline{G}$ is the argument of Λ_Δ corresponding to G. Let Algorithm 2 be the so-obtained algorithm. As stated next it is sound and complete.

Theorem 4. *If AF-Solver is sound and complete, for any goal element G of Δ, Algorithm 2 returns $SA_{u(\Delta)}(G)$ w.r.t. the updated ASAF $u(\Delta)$.*

4. Empirical Evaluation

We implemented a C++ prototype and compared the performance of: 1) the *incremental approach*, that is Algorithm 2 where AF-Solver is *μ-toksia* [35], the winner of the last ICCMA edition for the task DS-pr (i.e., computing the skeptical preferred acceptance of an argument of an AF); and 2) the *computation from scratch*, that is the computation of the skeptical preferred acceptance of the goal element w.r.t the updated ASAF by running AF-Solver (i.e., *μ-toksia*) directly on the AF for the updated ASAF.

Dataset. Although there are several benchmark generators and solvers for Dung's AFs [38], only a benchmark has been recently proposed for ASAFs [1]. Following [1], we generated a set of benchmark ASAFs by starting from AFs used as benchmark at ICCMA'19. Specifically, we use an AF dataset consisting of 326 AFs and, given a benchmark AF Λ, we generate an ASAF as follows: 30% of attacks in Λ are transformed into first-level supports; 12% (resp. 3%) of attacks in Λ are transformed into second-level supports towards a support (resp. an attack); 3% (resp. 2%) of attacks in Λ are transformed into third-level supports towards a support (resp. an attack); 12% (resp. 3%) of attacks in Λ are transformed into second-level attacks towards an attack (resp. a support); 2% (resp. 3%) of attacks in Σ are transformed into third-level attacks towards an attack (resp. a support); the remaining 30% of attacks in Λ are kept as first-level attacks of the

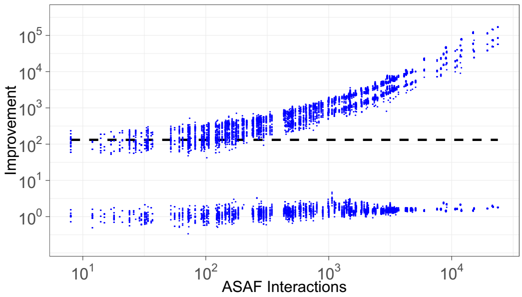

Figure 5. Improvement of the incremental approach over the computation from scratch (log scales). The dashed black line represents the median value.

resulting ASAF. This benchmark generation process aimed at preserving AFs' topology as much as possible. However, the process of generating ASAF benchmarks starting from AF benchmarks is challenging because we require specific amounts of different kind of attacks and supports, and we also need to check that the sub-graph induced by first-level supports is acyclic. Hence, to make it feasible, for each dataset, we generated an ASAF Δ if the number of arguments $|\mathbb{A}_\Delta|$ of the AF Λ_Δ for Δ does not exceed the number of arguments of the biggest AF in the original dataset. Therefore, starting from the AF dataset, we obtained an ASAF dataset consisting of 284 ASAFs $\Delta = \langle A, \Omega, \Gamma \rangle$ with a number of arguments $|A| \in [5, 10K]$ and a number of interactions $|\Omega \cup \Gamma| \in [8, 310K]$.

Methodology. For each ASAF Δ in the dataset, we consider a (randomly chosen) goal element and an update u selected among one of the possible 12 types (addition/deletion of an attack/support towards an argument/attack/support). Next, we compute the updated skeptical preferred acceptance of the goal element in the updated ASAF $u(\Delta)$ by calling Algorithm 2. Finally, we compute the *improvement* of Algorithm 2 over the computation from scratch as t_s/t_{A_2} where *i*) t_s is the time needed by the computation from scratch, and *ii*) t_{A_2} is the time needed by Algorithm 2. Thus, the improvement tells us how many times Algorithm 2 is faster than the computation from scratch. The experiments have been carried out on an Intel Core i7-3770K CPU 3.5GHz, 12GB RAM, running Ubuntu.

Results. Figure 5 reports the improvement versus the number of ASAF interactions (i.e., $|\Omega \cup \Gamma|$). Each data point refers to a run concerning an update and a goal. We also report the median of the improvement (dashed black line). Since μ-toksia ran into memory capacity saturation when computing the skeptical acceptance for $4,9\%$ of the AFs for the ASAFs in the dataset, we report the results for the remaining 244 ASAFs having number of arguments $|A| \in [5, 10K]$ and number of interactions $|\Omega \cup \Gamma| \in [8, 23.7K]$.

The results in Figure 5 show that, for a given goal and update, the improvement can be either very large or limited. This is due to the fact that either *i*) the alterable set is empty, and thus the algorithm immediately recognizes that acceptance status of the goal does not change after the update, or *ii)* the Proxy ASAF is built to compute the skeptical acceptance of the goal by invoking the external solver. Case *i*) occurs for 56% of the data points, and the average improvement in this case is 5836. The average improvement in the other case is 1.53, that is, the incremental computation takes 65% of the amount of time needed by the computation from scratch. In particular, although the size of the Proxy ASAF is 70.1% of that of the input ASAF on average, there is an overhead due to the construction of the Proxy ASAF that, to some extent, mitigates the benefit of the local computation on the smaller ASAF. Finally, the running time of Algorithm 2 is slightly more than that of the computation from scratch for only 3.8% of the data points. However, overall the incremental algorithm outperforms the computation from scratch, as confirmed by the median value of the improvements which is equal to 131 (the average is 3287, but is skewed by huge values of improvements in Figure 5).

5. Conclusions and Future Work

There has been an extensive body of work on managing changes in argumentation (a survey can be found in [28]). Besides the works mentioned in the introduction, other significant efforts coping with dynamics aspects of AFs include [10,14,21,23]. Similarly to what is done in this paper, some approaches focused on local computation in dynamic AFs [2,15,34,32] but with the aim of recomputing extensions. Recently, as discussed in Section 3.1, an algorithm for the incremental computation of an extension of dynamic ASAFs has been proposed in [1]. Moreover, an incremental approach to computing skeptical acceptance in Dung's frameworks has been proposed in [5], where the ideal extension is used for the computation and it is incrementally maintained. To the best of our knowledge, this is the first paper proposing an incremental technique for the computation of skeptical preferred acceptance in dynamic ASAFs. Due to the generality of ASAF, our technique can be also applied to restricted frameworks such as Argumentation Frameworks with Recursive Attacks (AFRAs) [13] and AFNs [36].

As future work we plan to investigate similar approaches for Recursive Argumentation Framework with Necessities (RAFN) [24], where a support may come also from a set of arguments, as well as extending our technique to deal with other semantics and considering the problem of enumerating extensions (as done for AFs [6]).

References

[1] G. Alfano, A. Cohen, S. Gottifredi, S. Greco, F. Parisi, and G. R. Simari. Dynamics in abstract argumentation frameworks with recursive attack and support relations. In *ECAI 2020 (to appear)*.

[2] G. Alfano, S. Greco, and F. Parisi. Efficient computation of extensions for dynamic abstract argumentation frameworks: An incremental approach. In *IJCAI*, pages 49–55, 2017.

[3] G. Alfano, S. Greco, and F. Parisi. Computing extensions of dynamic abstract argumentation frameworks with second-order attacks. In *IDEAS*, pages 183–192, 2018.

[4] G. Alfano, S. Greco, and F. Parisi. A meta-argumentation approach for the efficient computation of stable and preferred extensions in dynamic bipolar argumentation frameworks. *Intelligenza Artificiale*, 12(2):193–211, 2018.

[5] G. Alfano, S. Greco, and F. Parisi. An efficient algorithm for skeptical preferred acceptance in dynamic argumentation frameworks. In *IJCAI*, pages 18–24, 2019.

[6] G. Alfano, S. Greco, and F. Parisi. On scaling the enumeration of the preferred extensions of abstract argumentation frameworks. In *SAC*, pages 1147–1153, 2019.

[7] G. Alfano, S. Greco, F. Parisi, G. I. Simari, and G. R. Simari. Incremental computation of warranted arguments in dynamic defeasible argumentation: the rule addition case. In *SAC*, pages 911–917, 2018.

[8] G. Alfano, S. Greco, F. Parisi, G.I. Simari, and G.R. Simari. An incremental approach to structured argumentation over dynamic knowledge bases. In *KR*, pages 78–87, 2018.

[9] L. Amgoud, C. Cayrol, and M.-C. Lagasquie-Schiex. On the bipolarity in argumentation frameworks. In *NMR*, pages 1–9, 2004.

[10] L. Amgoud and S. Vesic. Revising option status in argument-based decision systems. *J. Log. Comp.*, 22(5):1019–1058, 2012.

[11] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, Henry Prakken, C. Reed, G. R. Simari, M. Thimm, and Serena Villata. Towards artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.

[12] P. Baroni, G. Boella, F. Cerutti, M. Giacomin, L. W. N. van der Torre, and S. Villata. On the input/output behavior of argumentation frameworks. *AI*, 217:144–197, 2014.

[13] P. Baroni, F. Cerutti, M. Giacomin, and G. Guida. AFRA: Argumentation Framework with Recursive Attacks. *IJAR*, 52(1):19–37, 2011.

[14] P. Baroni, M. Giacomin, and G. Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

[15] P. Baroni, M. Giacomin, and B. Liao. On topology-related properties of abstract argumentation semantics. A correction and extension to dynamics of argumentation systems: A division-based method. *AI*, 212:104–115, 2014.

[16] R. Baumann. Splitting an argumentation framework. In *LPNMR*, pages 40–53, 2011.

[17] T.J.M. Bench-Capon and P. E. Dunne. Argumentation in artificial intelligence. *AI*, 171:619 – 641, 2007.

[18] P. Bisquert, C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Characterizing change in abstract argumentation systems. In *Trends in Belief Revision and Argumentation Dynamics*, volume 48, pages 75–102. 2013.

[19] S. Bistarelli, F. Faloci, F. Santini, and C. Taticchi. Studying dynamics in argumentation with Rob. In *COMMA*, pages 451–452, 2018.

[20] G. Boella, D. M. Gabbay, L. W. N. van der Torre, and S. Villata. Support in abstract argumentation. In *COMMA*, pages 111–122, 2010.

[21] G. Boella, S. Kaci, and L. W. N. van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *ECSQARU*, pages 107–118, 2009.

[22] M. W. A. Caminada, W. Dvorák, and S. Vesic. Preferred semantics as socratic discussion. *J. of Log. and Comp.*, 26(4):1257–1292, 2016.

[23] C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Revision of an argumentation system. In *KR*, pages 124–134, 2008.

[24] C. Cayrol, J. Fandinno, L. Fariñas del Cerro, and M.-C. Lagasquie-Schiex. Structure-based semantics of argumentation frameworks with higher-order attacks and supports. In *COMMA*, pages 29–36, 2018.

[25] C. Cayrol and M.-C. Lagasquie-Schiex. Bipolarity in argumentation graphs: Towards a better understanding. *IJAR*, 54(7):876–899, 2013.

[26] G. Charwat, W. Dvorák, S. A. Gaggl, J. P. Wallner, and S. Woltran. Methods for solving reasoning problems in abstract argumentation - A survey. *AI*, 220:28–63, 2015.

[27] A. Cohen, S. Gottifredi, A. J. Garcia, and G. R. Simari. A survey of different approaches to support in argumentation systems. *The Knowl. Eng. Rev.*, 29(5):513–550, 2014.

[28] S. Doutre and J.-G. Mailly. Constraints and changes: A survey of abstract argumentation dynamics. *A & C*, 9(3):223–248, 2018.

[29] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *AI*, 77(2):321–358, 1995.

[30] P. E. Dunne and M. Wooldridge. Complexity of abstract argumentation. In *Argumentation in Artificial Intelligence*, pages 85–104. 2009.

[31] S. Gottifredi, A. Cohen, A. J. Garcia, and G. R. Simari. Characterizing acceptability semantics of argumentation frameworks with recursive attack and support relations. *AI*, 262:336–368, 2018.

[32] S. Greco and F. Parisi. Incremental computation of deterministic extensions for dynamic argumentation frameworks. In *JELIA*, pages 288–304, 2016.

[33] M. Kröll, R. Pichler, and S. Woltran. On the complexity of enumerating the extensions of abstract argumentation frameworks. In *IJCAI*, pages 1145–1152, 2017.

[34] B. Liao, L. J., and R. C. Koons. Dynamics of argumentation systems: A division-based method. *AI*, 175(11):1790–1814, 2011.

[35] A. Niskanen and M. Järvisalo. μ-toksia participating in ICCMA 2019. *Third ICCMA*, 2019.

[36] F. Nouioua and V. Risch. Argumentation frameworks with necessities. In *SUM*, pages 163–176, 2011.

[37] I. Rahwan and G. R. Simari. *Argumentation in Artificial Intelligence*. Springer, 2009.

[38] M. Thimm and S. Villata. The first international competition on computational models of argumentation: Results and analysis. *AI*, 252:267–294, 2017.

[39] S. Villata, G. Boella, D. M. Gabbay, and L. W. N. van der Torre. Modelling defeasible and prioritized support in bipolar argumentation. *AMAI*, 66(1-4):163–197, 2012.