On the Aggregation of Argumentation Frameworks: Operators and Postulates

Jérôme Delobelle   Sébastien Konieczny   Srdjan Vesic

CRIL, CNRS – Université d'Artois, Lens, France
{delobelle,konieczny,vesic}@cril.fr

Abstract

In this paper, we study the problem of aggregation of Dung's abstract argumentation frameworks. An argumentation framework allows the representation of conflictual agent's beliefs by using a set of arguments and interactions between them (*i.e.*, attack or non-attack). One argumentation framework per agent can be used to represent the beliefs of a group of agents. The aggregation process aims to represent the beliefs of this group by solving the potential conflicts between them. Some aggregation operators were defined, and more recently, some rationality properties for this process were introduced. In this work, we study the existing operators as well as some new ones, which we define in light of the proposed properties. We highlight the fact that existing operators do not satisfy a lot of properties. The conclusions are that on one hand none of the existing operators seem fully satisfactory, but on the other hand some of the properties proposed so far seem too demanding.

Key words: Argumentation, Aggregation

1 Introduction

Argumentation is based on the exchange and the evaluation of interacting arguments. In the last twenty years, it became one of the core subjects of Artificial Intelligence [4]. Indeed, argumentation can be used for modelling some aspects of reasoning, decision making, and dialogue; as such, it is now adopted as part of real-world applications from legal reasoning to sense-making in intelligence analysis. Among the existing formalisms, the work of Dung [19] on

* This is an extended version of the paper "Aggregation of Argumentation Frameworks" [18] written by the same authors and published in the 24th International Joint Conference on Artificial Intelligence, IJCAI 2015.

abstract argumentation is now recognized as a significant basis in the field. In Dung's abstract argumentation theory, the relation between arguments takes the form of conflicts, and the main goal is to extract a subset of arguments which can be jointly accepted w.r.t the way they interact. Dung's argumentation frameworks (AF) are modelled by directed graphs, where the nodes represent the abstract arguments and the edges represent the attacks between them. Several semantics (see [2] for an overview) were proposed to indicate which subsets of arguments, called extensions, are acceptable.

Studying dynamics of argumentation frameworks is a question of interest for researchers in the area. In recent years there have been a lot of contributions regarding the problem of evolution (revision) of argumentation frameworks [8,5,1,16,3,11,12]. However, there are few contributions on the aggregation of argumentation frameworks. This problem, illustrated in Fig. 1, is an important one for multi-agent systems: each agent may be associated with a different abstract argumentation framework on the same set of arguments (i.e. each agent may have different views on what constitutes a valid attack) that represents his beliefs. The problem is how to define a suitable representation of the beliefs of the group. This is an important question on its own, but also because this computed aggregated argumentation framework could be considered as the result of an *ideal* negotiation process, with respect to which the results of *practical* negotiation protocols (see e.g. the work by Bonzon and Maudet [6]) could be evaluated.

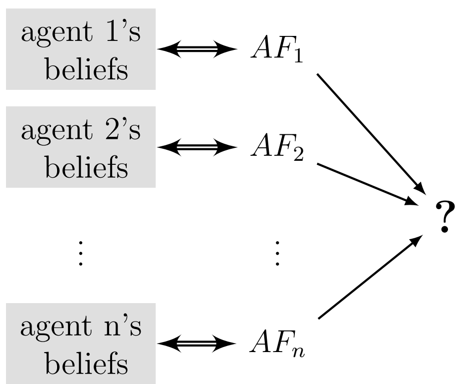

Fig. 1. Problem of aggregation of n AFs where each AF represents the beliefs of one agent

Contributions on this issue have been mainly proposals of particular aggregation methods [10,27,9,22]. We focus on the methods that comply with classical Dung's setting, meaning that they take as an input a classical abstract argumentation frameworks, and return as an output either an abstract argumentation framework (or more generally a set of abstract argumentation frameworks), or a set of extensions. Therefore, we do not study here proposals like the one by Gabbay and Rodrigues [22], since the output of their framework is an ordered set of arguments. Our work also differs from the work done by Caminada and Pigozzi [7] where the problem is how to find the correct labelling amongst the ones that correspond to a given argumentation framework while we consider several argumentation frameworks as input.

To aggregate several AFs, different questions arise like: *If a strict majority of agents think that an attack between two arguments exists, should this attack appear in the aggregation result? If an argument is accepted by all the agents, should this argument appear in the aggregation result?* Inspired by the social choice theory, Dunne et al. [21] formally define a number of axioms for argument aggregation. However, they do not check whether existing aggregation methods of AFs satisfy them.
This is what we propose in this work. We study existing methods from the literature in light of the proposed properties. We also propose an additional method based on WAFs (Weighted Argumentation Frameworks) [20,13] where one of the possible interpretations of the weights on the attacks that is mentioned is that they may represent the number of agents in a group that agree with this attack. We endorse this view, and check how to use these works to define aggregation methods. We assume that the set of arguments is identical for all the agents. This is common in many situations where all the agents receive the same information and accept each other's arguments. However, they often do not agree on whether an argument attacks another one [15]. Thus, the agents can have same arguments but different attack relations. In such a case, one of our operators could be useful. Of course, in some situations, the agents do not have the same knowledge and do not share all the arguments (e.g. because of privacy and/or strategical issues). This situation was studied by Coste-Marquis et al. [10]. They introduce a new attack relation between arguments, called the ignorance relation (an agent does not know whether there is an attack between two arguments). The next step of our work could be to propose aggregation operators in this more general setting.

The remainder of this paper is organised as follows. In the next section we provide the necessary background on argumentation frameworks. Section 3 is a reminder about weighted argumentation frameworks, and Section 4 recalls the proposals for aggregating argumentation frameworks. Section 5 studies the properties of these aggregation methods. In Section 6 we check the properties satisfied by the existing operators. In Section 7 we propose a new aggregation method, more precisely three variations, and study their properties. In Section 8 we sum up and discuss the obtained results, before concluding in Section 9.

2 Preliminaries

In this section, we briefly recall some key elements of abstract argumentation frameworks as proposed by Dung [19].

Definition 1 *An* **argumentation framework (AF)** *is a pair* $F = \langle A, R \rangle$ *with* A *a set of arguments and* R *a binary relation on* A*, i.e.* $R \subseteq A \times A$*, called the* **attack relation**. *A set of arguments* $\mathcal{S} \subseteq A$ *attacks an argument* $b \in A$*, if there exists* $a \in \mathcal{S}$*, such that* $(a, b) \in R$*. We use the notation* $Arg(F) = A$ *and* $Att(F) = R$*.*

The central question is to determine the sets of arguments that can be accepted together. Let us first introduce the notions of conflict-freeness and acceptability.

Definition 2 *Let* $F = \langle A, R \rangle$ *be an argumentation framework. A set of arguments* $\mathcal{S} \subseteq A$ *is* **conflict-free** *in* F *iff there exists no* $a, b \in \mathcal{S}$ *such that* $(a, b) \in R$*. An argument* $a \in A$ *is* **acceptable** *with respect to* $\mathcal{S}$ *iff for each* $b \in A$*, if* $(b, a) \in R$ *then* b *is attacked by* $\mathcal{S}$*.*

A set of arguments is **admissible** when it is conflict-free and each argument of the set is acceptable for this set. Given an argumentation framework one of the main questions is identifying sets $\mathcal{S}$ of arguments which can be accepted together (so called **extensions**). To find these extensions, several semantics have been proposed but, in this paper, we only focus on the standard semantics defined in [19]:

- $\mathcal{S}$ is a **complete** extension of F iff it is an admissible set and every argument which is acceptable with respect to $\mathcal{S}$ belongs to $\mathcal{S}$,
- $\mathcal{S}$ is a **preferred** extension of F iff it is a maximal (w.r.t. set inclusion $\subseteq$) admissible set of F,
- $\mathcal{S}$ is a **stable** extension of F iff it is a conflict-free set and it attacks all the arguments that do not belong to $\mathcal{S}$.
- $\mathcal{S}$ is a **grounded** extension of F iff $\mathcal{S}$ is the least fixed point of the characteristic function H of F defined by H: $2^A \rightarrow 2^A$ with H($\mathcal{S}$) = {a $\in A$: a is acceptable with respect to $\mathcal{S}$}.

We denote by $\mathcal{E}_\sigma$(F) the set of extensions of F for the semantics $\sigma \in$ {**comp**(lete), **pref**(erred), **sta**(ble), **gr**(ounded)}.

We can now define the acceptability status of each subset of arguments. Given a semantics σ, an argument a is **skeptically** accepted iff there exists at least one extension and a belongs to all extensions. An argument is **credulously** accepted iff it belongs to at least one extension. We denote by $sa_\sigma(F)$ (resp. $ca_\sigma(F)$) the set of skeptically (resp. credulously) accepted arguments in F.

3 Weighted Argumentation Frameworks

Let us now turn to an extension of Dung's framework where a weight on each attack is added [20,13,14]. Several interpretations of weights on attacks exist [20]: a measure of the extent of inconsistency between pairs of arguments (the higher the weight on the attack between two arguments, the greater the inconsistency between them), the relative strength of the attack (higher weight denotes a stronger attack), or the number of agents that support the attack.

Definition 3 **[13]** *A* Weighted Argumentation Framework (WAF) *is a triple* $\mathrm{WF} = \langle A, R, w \rangle$ *where* $\langle A, R \rangle$ *is a Dung abstract argumentation framework, and* $w : A \times A \to \mathbb{N}$ *is a function assigning a natural number to each attack (i.e.* $w(a, b) > 0$ *iff* $(a, b) \in R$*), and a null value otherwise* ($w(a, b) = 0$ *iff* $(a, b) \notin R$*).*

Example 1 *Let* $WF = \langle A, R, w \rangle$ *be a weighted argumentation framework with* $A = \{a, b, c, d, e, f\}$*,* $R = \{(c, a), (c, b), (d, c), (d, e), (e, d), (e, f)\}$ *and w:* $(c, a) \to 3$*,* $(c, b) \to 5$*,* $(d, c) \to 5$*,* $(d, e) \to 1$*,* $(e, d) \to 2$*,* $(e, f) \to 5$*.*

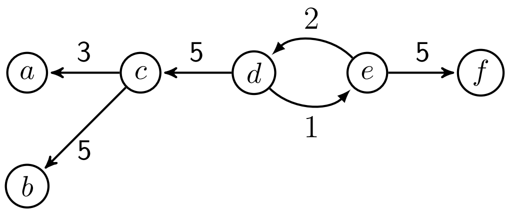

Fig. 2. A weighted argumentation framework

Let us denote as $\overline{\mathrm{WF}}$ the standard argumentation framework obtained from a weighted argumentation framework WF by "forgetting" the weights, i.e. if $\mathrm{WF} = \langle A, R, w \rangle$ then $\overline{\mathrm{WF}} = \langle A, R \rangle$.
We later show how to use WAFs to aggregate several AFs. For the moment let us recall how one can use them for the "relaxing" of extensions and for selecting the best amongst several extensions.

3.1 Relaxed Extensions

Initially WAFs were introduced with the idea to ensure non-empty extensions [20]. Roughly, the goal is to delete some attacks in order to obtain a non-empty set of extensions.

Definition 4 **[14]** *We say that* $\oplus$ *is an aggregation function if for every* $n \in \mathbb{N}$*,* $\oplus$ *is a mapping from* $\mathbb{N}^n$ *to* $\mathbb{N}$ *such that:*

- *if* $x_i \geq x'_i$, *then* $\oplus(x_1, \dots, x_i, \dots, x_n) \geq \oplus(x_1, \dots, x'_i, \dots, x_n)$ **(non-decreasingness)**
- $\oplus(x_1, \dots, x_n) = 0$ *iff for every* i, $x_i = 0$ **(minimality)**
- $\oplus(x) = x$ **(identity)**

Thus, given a natural number β, an aggregation function is used to select the set of attacks such that the aggregation of their weight is smaller than β.

Definition 5 [14] *Let* $\mathrm{WF} = \langle A, R, w\rangle$ *be a weighted argumentation framework,* σ *be a semantics, and* $\otimes$ *be an aggregation function. The aggregation of the weights of the attacks in a set* $S \subseteq R$ *is* $w_{\otimes}(S, w) = \otimes_{(a,b)\in S} w(a, b)$*. The function Sub(R,w,β) returns the set of subsets of R whose total aggregated weight does not exceed* β*:* $Sub(R, w, \beta) = \{S \mid S \subseteq R \text{ and } w_{\otimes}(S, w) \leq \beta\}$*. The set of* $\sigma^{\beta}_{\otimes}$*-extensions of* WF*, denoted* $\mathcal{E}^{\otimes,\beta}_{\sigma}(\mathrm{WF})$*, is defined as:*

$$\mathcal{E}^{\otimes,\beta}_{\sigma}(\mathrm{WF}) = \{E \in \mathcal{E}_{\sigma}(\langle A, R\backslash S\rangle) \mid S \in Sub(R, w, \beta)\}$$

Example 1 (cont.) *Let us compute the set of* $\sigma^{\beta}_{\otimes}$*-extensions of the WAF illustrated in Fig. 2 for different values of* β*. We choose the sum as the aggregation function* ($\otimes = \Sigma$) *and the preferred and grounded semantics* ($\sigma = \{pref, gr\}$)*.*

β	$\mathcal{E}^{\Sigma,\beta}_{pref}(\mathrm{WF})$	$\mathcal{E}^{\Sigma,\beta}_{gr}(\mathrm{WF})$
0	$\{\{c, e\}, \{a, b, d, f\}\}$	$\{\emptyset\}$
1	$\{\{c, e\}, \{a, b, d, f\}\}$	$\{\emptyset, \{c, e\}\}$
2	$\{\{c, e\}, \{a, b, d, f\}\}$	$\{\emptyset, \{c, e\}, \{a, b, d, f\}\}$
3	$\{\{c, e\}, \{a, c, e\}, \{a, b, d, f\}, \{a, b, d, e\}\}$	$\{\emptyset, \{a\}, \{c, e\}, \{a, b, d, f\}, \{a, b, d, e\}\}$

Table 1
$\mathcal{E}^{\otimes,\beta}_{\sigma}(\mathrm{WF})$ for $\beta \in \{0, 1, 2, 3\}$

We can see that there exists no non-empty grounded extension associated with the classical argumentation framework $\overline{\mathrm{WF}}$ *:* $\mathcal{E}_{gr}(\overline{\mathrm{WF}}) = \{\emptyset\}$*. However, in using the* relaxed extensions*, new extensions are generated. For example, when* $\beta = 2$*, it is possible to remove either the attack from d to e (in this case we obtain the extension* $\{c, e\}$*) or the attack from e to d (we obtain the extension* $\{a, b, d, f\}$*). Please note that these two attacks are the only ones that are removable because they have a weight lesser or equal to* β*. So we add these two extensions to previous ones* ($\beta < 2$) *to obtain* $\mathcal{E}^{\Sigma,2}_{gr}(\mathrm{WF}) = \{\emptyset, \{c, e\}, \{a, b, d, f\}\}$*.*

In contrast to what happens in Dung's setting, several grounded extensions may exist when the relaxing of extensions is considered. Furthermore, it may be the case that the empty set belongs to a set of relaxed extensions. This situation can be problematic because it trivializes skeptical inference relation, so the next definition removes it from the set of extensions. Thus, the most interesting value of β is the smallest one that leads to a non-empty extension (for the semantics under consideration):

Definition 6 *Given a weighted argumentation framework* WF*, a semantics* σ*, and an aggregation function* $\otimes$*, the set of* $\sigma_{\otimes}$*-extensions of WF, denoted by* $\mathcal{E}_{\sigma}^{\otimes}(\mathrm{WF})$ *is defined as* $\mathcal{E}_{\sigma}^{\otimes}(\mathrm{WF}) = \underline{\mathcal{E}_{\sigma}^{\otimes,\beta}(\mathrm{WF})}$ *where*[1] *:*

- $\mathcal{E}_{\sigma}^{\otimes,\beta}(\mathrm{WF})$ *is non-trivial*[2] *,*
- *there is no* $\beta' < \beta$ *s.t.* $\mathcal{E}_{\sigma}^{\otimes,\beta'}(\mathrm{WF})$ *is non-trivial,*
- *for a set* $\mathcal{E}$ *of extensions,* $\underline{\mathcal{E}} = \mathcal{E} \setminus \{\emptyset\}$*.*

Example 1 (cont.) *There exists no change for the preferred semantics which is already non-trivial when* $\beta = 0$*, so* $\mathcal{E}_{pref}^{\Sigma}(\mathrm{WF}) = \{\{c, e\}, \{a, b, d, f\}\}$*. However, for the grounded semantics, the case* $\beta = 0$ *does not respect the first condition concerning the non-triviality. So the smallest value of* β *that leads to a non-trivial set of extensions is when* $\beta = 1$*. So, after removing the empty set, we obtain* $\mathcal{E}_{gr}^{\Sigma}(\mathrm{WF}) = \{\{c, e\}\}$*.*

3.2 Best Extensions

In general, an argumentation framework may admit a large number of extensions for some semantics. Within the WAF setting, it is possible to take advantage of the available weights, in order to select the "best" extensions. In the paper by Coste-Marquis et al. [13] this selection process goes through a comparison of the extensions' scores, expressing intuitively how good they are.

Definition 7 *Let* $\mathrm{WF} = \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* E *and* F *be two extensions of* $\overline{\mathrm{WF}}$ *for a given semantics* σ *and* $\oplus$ *be an aggregation function. The* ***$\oplus$-attack*** *from* E *on* F *is:* $S_{\oplus}(E \to F) = \oplus_{a \in E, b \in F} w(a, b)$*. Then* $E >_{\oplus} F$ *iff* $S_{\oplus}(E \to F) > S_{\oplus}(F \to E)$*.*

Let us introduce an additional parameter to the best function, that will be useful for the operators we define in Section 7.

Definition 8 *Let* $\mathrm{WF} = \langle A, R, w\rangle$ *be a weighted argumentation framework and* $X \subseteq 2^{A}$*. Let* $\oplus$ *be an aggregation function. Then :*

- $best_1^{\oplus}(X, \mathrm{WF}) = \{E \in X : \nexists E' \in X, E' >_{\oplus} E\}$
- $best_2^{\oplus}(X, \mathrm{WF}) = argmax_{E \in X} |\{E' \in X : E >_{\oplus} E'\}|$
- $best_3^{\oplus}(X, \mathrm{WF}) = argmax_{E \in X} |\{E' \in X : E >_{\oplus} E'\}| - |\{E' \in X : E' >_{\oplus} E\}|$
- $best_4^{\oplus}(X, \mathrm{WF}) = argmax_{E \in X} KS_{\oplus}(E)$,
  where $KS_{\oplus}(E) = min_{E' \in X, E' \neq E}(S_{\oplus}(E \to E'))$

[1] We add a third condition compared to the original definition [14], in order to directly encode the non-trivial skeptical inference [14].

[2] A set of extensions is non-trivial if it has at least one non-empty extension.

The first rule selects the extensions which are not beaten by any other extension. This method can give an empty set as the answer, which is not the case with the other three methods. The second approach consists in counting how many extensions are defeated by a given extension. The third rule is similar to the Copeland voting rule [25] because the score of an extension E is the difference between the number of extensions beaten by E and the number of extensions that beat E. The fourth method, similar to the Kramer-Simpson voting rule [24,26], selects the extensions maximising the KS score, which is the minimal value of $\oplus$-attack from E to E', when E' ranges over all extensions of $\overline{\mathrm{WF}}$.

Four natural ways to define the ordering $>_\oplus$ are proposed in [13] :

Definition 9 *Let WF be a weighted AF. Let $\oplus$ be an aggregation function. Then, $\forall i \in \{1, 2, 3, 4\}$, $best_i^{\sigma,\oplus}(\mathrm{WF}) = best_i^{\oplus}(\mathcal{E}_\sigma(\overline{\mathrm{WF}}), \mathrm{WF})$.*

Example 2 [13] *Let us compute the best extensions of the weighted argumentation framework* WF *illustrated in Fig. 3.*

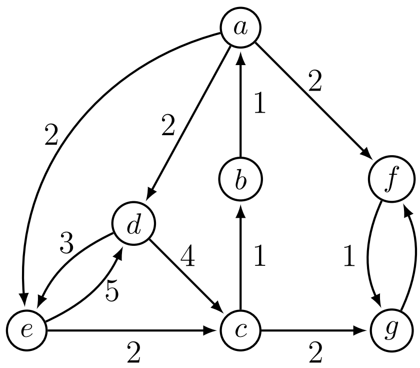

Fig. 3. The digraph of WF

$\overline{\mathrm{WF}}$ *has five preferred extensions : $\mathcal{E}_{pref}(\overline{\mathrm{WF}}) = \{\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_4, \mathcal{E}_5\}$ with $\mathcal{E}_1 = \{a, c\}$, $\mathcal{E}_2 = \{b, d, f\}$, $\mathcal{E}_3 = \{b, e, f\}$, $\mathcal{E}_4 = \{b, d, g\}$ and $\mathcal{E}_5 = \{b, e, g\}$. To compare these extensions, we use the $\oplus$-attack (see Definition 7), focusing on two aggregation functions $\oplus = \{\Sigma, max\}$:*

Then, in using the different methods best, the following extensions are selected:

- $best_1^{pref,\Sigma}(\mathrm{WF}) = \{\mathcal{E}_1\}$
- $best_2^{pref,\Sigma}(\mathrm{WF}) = \{\mathcal{E}_5\}$
- $best_3^{pref,\Sigma}(\mathrm{WF}) = \{\mathcal{E}_1, \mathcal{E}_5\}$
- $best_4^{pref,\Sigma}(\mathrm{WF}) = \{\mathcal{E}_1\}$
- $best_1^{pref,max}(\mathrm{WF}) = \{\mathcal{E}_5\}$
- $best_2^{pref,max}(\mathrm{WF}) = \{\mathcal{E}_5\}$
- $best_3^{pref,max}(\mathrm{WF}) = \{\mathcal{E}_5\}$
- $best_4^{pref,max}(\mathrm{WF}) = \{\mathcal{E}_4\}$

Σ	$\mathcal{E}_1$	$\mathcal{E}_2$	$\mathcal{E}_3$	$\mathcal{E}_4$	$\mathcal{E}_5$
$\mathcal{E}_1$	$\times$	5	5	5	5
$\mathcal{E}_2$	5	$\times$	3	1	4
$\mathcal{E}_3$	3	5	$\times$	6	1
$\mathcal{E}_4$	5	4	7	$\times$	3
$\mathcal{E}_5$	3	9	4	5	$\times$

max	$\mathcal{E}_1$	$\mathcal{E}_2$	$\mathcal{E}_3$	$\mathcal{E}_4$	$\mathcal{E}_5$
$\mathcal{E}_1$	$\times$	2	2	2	2
$\mathcal{E}_2$	4	$\times$	3	1	3
$\mathcal{E}_3$	2	5	$\times$	5	1
$\mathcal{E}_4$	4	4	4	$\times$	3
$\mathcal{E}_5$	2	5	4	5	$\times$

Fig. 4. Comparison of extensions from $\overline{\mathrm{WF}}$ using the $\oplus$-attack ($\oplus \in \{\Sigma, max\}$). Each number in the table corresponds to the score obtained from the attacks from an extension in a row on an extension in the corresponding column.

4 Aggregation Operators

Some merging operators for argumentation frameworks, inspired by propositional logic merging operators [23], have been defined in the literature. From now on, we make the hypothesis that all the argumentation frameworks are defined on the same set of arguments $\mathcal{X}$. We use the notation $\mathbb{AF}$ for the set of all argumentation frameworks that can be defined from the set of arguments $\mathcal{X}$ used by the agents. A profile of n AFs is denoted by $\hat{F} = (F_1, \dots, F_n)$ which is a tuple in $\mathbb{AF}^n$.

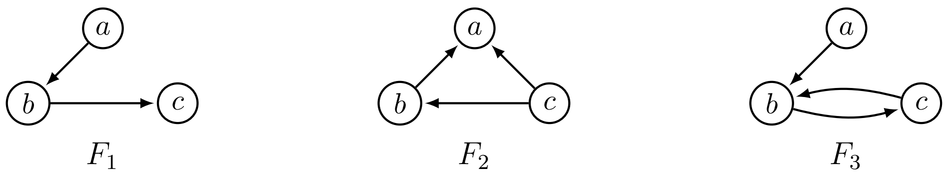

Fig. 5. A profile of three AFs

Coste-Marquis et al. [10] were the first to consider the problem of aggregating several argumentation frameworks. Their method is a two-phase process. In the first phase, each argumentation framework (which have not necessarily the same set of arguments) is converted into a partial argumentation framework which is an extension of Dung's framework where a new relation, called the ignorance relation (the agent does not know whether there is an attack or not between two arguments), is added. However, as we said in the introduction of this section, we only consider the AFs defined from the set of arguments, so we can directly go to the second step.

During the second step, where all the argumentation frameworks have the same set of arguments, a distance-based method is used to aggregate them. Indeed, the result of the aggregation is represented by one or several AFs which are as close as possible to the given profile of AFs.

Definition 10 *Let $\mathcal{D}$ be a distance between AFs, $\hat{F}$ be a profile and $\oplus$ be an aggregation function. The merging operator $\Delta^{\oplus}_{\mathcal{D}}$ is defined as :*

$$\Delta^{\oplus}_{\mathcal{D}}(\hat{F}) = \{F \in \mathbb{AF} : F \text{ minimizes } \oplus_{i=1}^{n} \mathcal{D}(F, F_i)\}$$

The distance $\mathcal{D}$ used by Coste-Marquis et al. [10] as an example is the edit distance (de), which is in our case equivalent to the cardinality of the symmetrical difference between two attack relations. Typical examples of aggregation functions are *sum* (Σ) and *leximax* [3]. In what follows, we focus on these two functions.

Example 3 *Let us compute the result of the aggregation of the profile represented in Fig. 5. From all of the possible AFs with three arguments that can be computed, we select the closest AFs to the profile in the input (see Fig. 6). With*

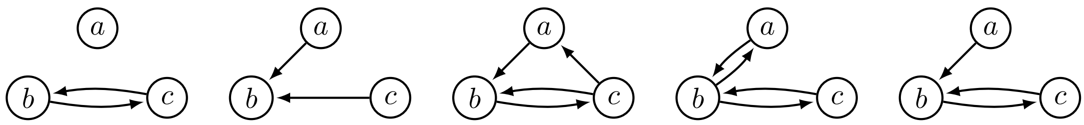

F'_1 F'_2 F'_3 F'_4 F'_5

Fig. 6. Set of possible AFs representing the result of aggregation $\Delta^{\oplus}_{\mathcal{D}}$

the sum as aggregation function and the edit distance (de) as distance between AFs, the only AF to obtain the minimal score is F'_5 ($\sum_{i=1}^{3} de(F'_5, F_i) = 5$), so $\Delta^{\Sigma}_{de}(\langle F_1, F_2, F_3\rangle) = \{F'_5\}$. However, with the leximax, there are four AFs with the minimal score : F'_1, F'_2, F'_3 and F'_4 (for the case of F'_1, we obtain $\mathcal{L}eximax_{i=1}^{3} de(F'_1, F_i) = (3, 2, 1)$), so $\Delta^{leximax}_{de}(\langle F_1, F_2, F_3\rangle) = \{F'_1, F'_2, F'_3, F'_4\}$.

Inspired by voting methods, other aggregation operators have been defined by Tohmé et al. [27]. In particular, they propose a *qualified voting* method where an attack ($a \to b$) appears in the unique resulting AF iff it appears more than its opposite attack ($b \to a$) but also more than the absence of attacks ($a \not\to b$ and $b \not\to a$) and a subset of agents U believes that this attack exists:

Definition 11 *Let $\hat{F}$ be a profile and $U \subset \{1, \ldots, n\}$ be a subset of agents. Let us note :*

- $S_{a\to b} = |\{F_i : (a, b) \in Att(F_i)\}|$
- $S_{b\to a} = |\{F_i : (b, a) \in Att(F_i)\}|$
- $S_{a\not\to b} = |\{F_i : (a, b) \notin Att(F_i) \text{ and } (b, a) \notin Att(F_i)\}|$

[3] When applied to a vector of n real numbers, the leximax function gives the list of those numbers sorted in a decreasing way. Such lists are compared w.r.t. the lexicographic ordering induced by the standard ordering on real numbers.

Qualified voting is defined as $QV(\hat{F}) = \langle A, R\rangle$*, where* A *is the set of arguments used by agents, and* $R = \{(a,b) \mid a,b \in A \text{ and } S_{a\to b} > max(S_{b\to a}, S_{a\nrightarrow b}) \text{ and } U \subseteq \{i : (a,b) \in Att(F_i)\}\}$.

This semantics has been defined in order to satisfy some properties inspired from social choice theory, in particular the non-dictatorship property (we will see the other properties in Section 5). To satisfy it, Tohmé et al. suppose that the set U should contain at least two agents ($|U| \geq 2$). In addition, it is important to note that the set U cannot contain all the agents.

Example 3 (cont.) *From the profile illustrated in Fig. 5, with* $U = \{1, 3\}$*, the result of the aggregation method QV is the AF depicted in Fig. 7 which contains only the attack from* a *to* b*.*

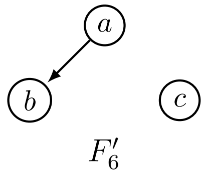

Fig. 7. Result returned by the aggregation method QV

Indeed, all the agents in U think that this attack exists ($(a,b) \in Att(F_1)$ *and* $(a,b) \in Att(F_3)$) *and* $S_{a\to b} = 2$, $S_{b\to a} = 1$ *and* $S_{a\nrightarrow b} = 0$*, which implies* $S_{a\to b} = 2 > 1 = max(S_{b\to a}, S_{a\nrightarrow b})$*. So* $QV(\langle F_1, F_2, F_3\rangle) = F'_6$*.*

5 Properties of Aggregation Function

Dunne et al. [21] propose some rationality properties for characterizing aggregation of a set of argumentation frameworks, based on translations of properties coming from social choice theory to the argumentation setting.

Recall that we denote by $\mathbb{AF}$ the set of all argumentation frameworks defined from a (finite) set of arguments and that we suppose that all agents have the same set of arguments. We denote by $\mathcal{N}$ the set of all agents. An aggregation function γ is defined by $\gamma : \mathbb{AF}^n \to \mathbb{AF}$. Unless stated explicitly all the properties are defined $\forall \hat{F} \in \mathbb{AF}^n$.

Anonymity[21] The aggregation function γ is **anonymous** if it produces the same argumentation framework for all permutations $\Pi(\hat{F})$ of the input.

$$\forall \hat{F}' \in \Pi(\hat{F}) : \gamma(\hat{F}) = \gamma(\hat{F}') \qquad \textbf{(ANON)}$$

Non-Triviality[21] An argumentation framework is non-trivial, for a semantics σ, if it has at least one non-empty extension: $|\mathcal{E}_\sigma(F)| \geqslant 1$ and $\mathcal{E}_\sigma(F) \neq$

$\{\emptyset\}$. Let us note $\mathbb{AF}_{NT_\sigma}$ the set of non-trivial (for the semantics σ) argumentation frameworks. The aggregation function γ is **σ-strongly non-trivial** if the ouput is always non-trivial:

$$\gamma(\hat{F}) \in \mathbb{AF}_{NT_\sigma} \qquad (\sigma\text{-}\mathbf{SNT})$$

The aggregation function γ is **σ-weakly non-trivial** if, when all the input frameworks are non-trivial, then the output framework is also non-trivial:

$$\forall \hat{F} \in \mathbb{AF}^n_{NT_\sigma} : \gamma(\hat{F}) \in \mathbb{AF}_{NT_\sigma} \qquad (\sigma\text{-}\mathbf{WNT})$$

Decisiveness[21] An argumentation framework is decisive, for a semantics σ, if it has exactly one non-empty extension: $|\mathcal{E}_\sigma(F)| = 1$ and $\mathcal{E}_\sigma(F) \neq \{\emptyset\}$. Let us note $\mathbb{AF}_{D_\sigma}$ the set of decisive (for the semantics σ) argumentation frameworks. The aggregation function γ is **σ-strongly decisive** if the output is always decisive:

$$\gamma(\hat{F}) \in \mathbb{AF}_{D_\sigma} \qquad (\sigma\text{-}\mathbf{SD})$$

The aggregation function γ is **σ-weakly decisive** if, when all the input frameworks are decisive, then the output framework is also decisive:

$$\forall \hat{F} \in \mathbb{AF}^n_{D_\sigma} : \gamma(\hat{F}) \in \mathbb{AF}_{D_\sigma} \qquad (\sigma\text{-}\mathbf{WD})$$

Unanimity[21] If all agents are unanimous with respect to some aspect of the domain (extensions, attacks, ...), for a semantics σ, then this unanimity should be reflected in the social outcome.

• **Unanimous attack** concerns attacks between arguments:

$$\bigcap_{k=1}^{n} Att(F_k) \subseteq Att(\gamma(\hat{F})) \qquad (\mathbf{A\text{-}U})$$

• **σ-unanimity** concerns extensions:

$$\bigcap_{k=1}^{n} \mathcal{E}_\sigma(F_k) \subseteq \mathcal{E}_\sigma(\gamma(\hat{F})) \qquad (\sigma\text{-}\mathbf{U})$$

• **ca_σ-unanimity** concerns credulous inference:

$$\bigcap_{k=1}^{n} ca_\sigma(F_k) \subseteq ca_\sigma(\gamma(\hat{F})) \qquad (ca_\sigma\text{-}\mathbf{U})$$

• **sa_σ-unanimity** concerns skeptical inference:

$$\bigcap_{k=1}^{n} sa_\sigma(F_k) \subseteq sa_\sigma(\gamma(\hat{F})) \qquad (sa_\sigma\text{-}\mathbf{U})$$

Majority[21] If a strict majority of agents agree on something, then this should be reflected in the social outcome:

• **Majority attack** concerns attacks between arguments:

$$(|\{F_i : a \in Att(F_i)\}| > \tfrac{n}{2}) \Rightarrow a \in Att(\gamma(\hat{F})) \qquad (\mathbf{A\text{-}MAJ})$$

• **σ-majority** concerns extensions:

$$(|\{F_i : S \in \mathcal{E}_\sigma(F_i)\}| > \tfrac{n}{2}) \Rightarrow S \in \mathcal{E}_\sigma(\gamma(\hat{F})) \qquad (\sigma\text{-}\mathbf{MAJ})$$

• **ca_σ-majority** concerns credulous inference:

$$(|\{F_i : x \in ca_\sigma(F_i)\}| > \tfrac{n}{2}) \Rightarrow x \in ca_\sigma(\gamma(\hat{F})) \qquad (ca_\sigma\text{-}\mathbf{MAJ})$$

• **sa_σ-majority** concerns skeptical inference:

$$(|\{F_i : x \in sa_\sigma(F_i)\}| > \tfrac{n}{2}) \Rightarrow x \in sa_\sigma(\gamma(\hat{F})) \qquad (sa_\sigma\text{-}\mathbf{MAJ})$$

Closure[21] The aggregation function must not invent entities that do not exist in the input.

• **Closure** says that the resulting AF must match exactly one AF in the input:

$$\exists i \in \mathcal{N} : Att(\gamma(\hat{F})) = Att(F_i) \qquad \textbf{(CLO)}$$

• **Attack closure** says that if one attack is in the resulting AF, this attack must be present in at least one AF in the input:

$$Att(\gamma(\hat{F})) \subseteq Att(F_1) \ \cup \ \ldots \ \cup \ Att(F_n) \qquad \textbf{(A-C)}$$

• σ**-closure** is related to extensions:

$$\forall S \in \mathcal{E}_\sigma(\gamma(\hat{F})) : S \in \bigcup_{k=1}^{n} \mathcal{E}_\sigma(F_k) \qquad (\sigma\textbf{-C})$$

• ca_σ**-closure** is related to credulous inference:

$$\forall x \in ca_\sigma(\gamma(\hat{F})) : x \in \bigcup_{k=1}^{n} ca_\sigma(F_k) \qquad (ca_\sigma\textbf{-C})$$

• sa_σ**-closure** is related to skeptical inference:

$$\forall x \in sa_\sigma(\gamma(\hat{F})) : x \in \bigcup_{k=1}^{n} sa_\sigma(F_k) \qquad (sa_\sigma\textbf{-C})$$

Tohmé et al. [27] propose some properties inspired from social choice theory for characterizing good aggregation operators. A property, called *Pareto condition*, is first introduced and is exactly Unanimous attack (A-U). A non-dictatorship property is then introduced and is satisfied by all reasonable aggregation operators. We give below the two other properties that they propose, that are translations of meaningful social choice theory properties:

Positive responsiveness [27] This property says that increasing the number of agents that have an attack should not decrease the chance for that attack to appear in the social outcome:

(PR) Let $\hat{F}$ and $\hat{G}$ be two profiles of $\mathbb{AF}^n$. If $\{F_i \in \hat{F} : (a,b) \in Att(F_i)\} \subseteq \{G_i \in \hat{G} : (a,b) \in Att(G_i)\}$ and $(a,b) \in Att(\gamma(\hat{F}))$, then $(a,b) \in Att(\gamma(\hat{G}))$

Independence of irrelevant alternatives [27] Deciding whether an attack holds or not should be concerned only with the attacks between these two arguments in the input profile.

(IIA) Let $\hat{F}$ and $\hat{G}$ be two profiles of $\mathbb{AF}^n$. If ($\forall i$ $(a,b) \in Att(F_i)$ iff $(a,b) \in Att(G_i)$), then ($(a,b) \in Att(\gamma(\hat{F}))$ iff $(a,b) \in Att(\gamma(\hat{G}))$)

This is not mentioned in [27], but one can easily show that PR implies IIA:

Proposition 1 *Positive responsiveness implies Independence of Irrelevance Alternatives.*

Proof. *The property PR says that for two profiles $\hat{F}$ and $\hat{G}$, if $\{F_i \in \hat{F} :$*

$(a,b) \in Att(F_i)\} \subseteq \{G_i \in \hat{G} : (a,b) \in Att(G_i)\}$, and $(a,b) \in Att(\gamma(\hat{F}))$, then $(a,b) \in Att(\gamma(\hat{G}))$. This is also true in the particular case when $\{F_i \in \hat{F} : (a,b) \in Att(F_i)\} = \{G_i \in \hat{G} : (a,b) \in Att(G_i)\}$. To prove IIA, let us suppose that $\forall i$ $(a,b) \in Att(F_i)$ iff $(a,b) \in Att(G_i)$. From PR we obtain that: if $(a,b) \in Att(\gamma(\hat{F}))$ then $(a,b) \in Att(\gamma(\hat{G}))$. Since $\{F_i \in \hat{F} : (a,b) \in Att(F_i)\} = \{G_i \in \hat{G} : (a,b) \in Att(G_i)\}$, the implication is also true in the other direction: $(a,b) \in Att(\gamma(\hat{G})) \Rightarrow (a,b) \in Att(\gamma(\hat{F}))$. Thus, IIA is satisfied. ■

To this set of properties from the literature we want to add the following intuitive properties.

Identity If all the AFs in the input coincide and are non-trivial, the result of merging should be identical to this AF.

• **Identity attack** on the attacks :

$$Att(\gamma(F, \ldots, F)) = Att(F) \qquad \textbf{(A-ID)}$$

• **σ-identity** on the extensions :

$$\forall F \in \mathbb{AF}_{NT_\sigma} : \mathcal{E}_\sigma(\gamma(F, \ldots, F)) = \mathcal{E}_\sigma(F) \qquad \textbf{(}\sigma\textbf{-ID)}$$

• **ca_σ-identity** on the credulous inference :

$$\forall F \in \mathbb{AF}_{NT_\sigma} : ca_\sigma(\gamma(F, \ldots, F)) = ca_\sigma(F) \qquad \textbf{(}ca_\sigma\textbf{-ID)}$$

• **sa_σ-identity** on the skeptical inference :

$$\forall F \in \mathbb{AF}_{NT_\sigma} : sa_\sigma(\gamma(F, \ldots, F)) = sa_\sigma(F) \qquad \textbf{(}sa_\sigma\textbf{-ID)}$$

These four intuitive properties are particular cases of the property of unanimity. In the case where we aggregate a single AF (which represents the beliefs of one agent for example) then this property implies that the operator should not change this AF, that, as we will see, is not ensured by all existing aggregation methods.

Proposition 2 *For every semantics σ such that there exists an argumentation framework F which is not decisive, the properties σ-strong decisiveness and σ-Identity are not compatible.*

Proof. *Let F be a non-decisive argumentation framework for a given semantics σ ($F \notin \mathbb{AF}_{D_\sigma}$) and $\hat{F} = (F, \ldots, F)$ be a profile. According to the property σ-Identity, the result of aggregation, and specifically its set of extensions, should be exactly the same as the input i.e. non-decisive ($\gamma(F, \ldots, F) \notin \mathbb{AF}_{D_\sigma}$). Conversely, the property of σ-strong decisiveness says that the result should be decisive ($\gamma(F, \ldots, F) \in \mathbb{AF}_{D_\sigma}$).* ■

6 Properties of Existing Operators

Let us first check which properties are satisfied by the aggregation operators of Coste-Marquis et al. [10]. Recall that the properties introduced by Dunne et al. are defined for a unique AF as an output, whereas the merging operators

may have several AFs as an output. We generalize the properties from the previous section as follows: instead of asking that a property holds for *the* AF at the output, we ask that the same property is satisfied by *all* AFs in the output. By following this idea, most of the properties can be generalized in a straightforward way. We generalize Positive Responsiveness as follows:
(PR) Let $\hat{F}$ and $\hat{G}$ be two profiles of $\mathbb{AF}^n$. If $\{F_i \in \hat{F} : (a,b) \in Att(F_i)\} \subseteq \{G_i \in \hat{G} : (a,b) \in Att(G_i)\}$, and $\forall F_j \in \gamma(\hat{F})$ we have $(a,b) \in Att(F_j)$, then $\forall G_k \in \gamma(\hat{G})$ we have $(a,b) \in Att(G_k)$.

There exists no straightforward way to generalize the definition of IIA to the case when several AFs are allowed in the output. That is why we do not consider it in the remainder of the paper.

Recall that the aggregation operators proposed in [10] have two parameters: a distance between AFs, and an aggregation function. Thus, we will use the edit distance (de) combined with the sum Σ (Proposition 3) or the leximax (Proposition 4) as aggregation function.

Proposition 3 *Let $\sigma \in \{comp, pref, sta, gr\}$ be a semantics. Δ^{Σ}_{de} satisfies Anonymity (ANON), the properties of Identity (A-ID, σ-ID, ca_σ-ID, sa_σ-ID), Unanimous attack (A-U), Majority attack (A-MAJ), Attack closure (A-C) and Positive responsiveness (PR). The other properties are not satisfied.*

Proof. *For Anonymity, note that Δ^{Σ}_{de} does not take into consideration the order of merging of the AFs.*
For the properties of Identity, the proof follows directly from the observation that $de(F, F') = 0$ if and only if $F = F'$. Thus, the unique argumentation framework minimizing the sum of distances is the one present at the input.
For the properties of attacks, a link has been established by Coste-Marquis and al. [10] between Δ^{Σ}_{de} and the majority graph where an attack is in this graph iff this attack is in a strict majority of AFs. The three properties (A-U, A-MAJ and A-C) are thus satisfied.
For PR, suppose that $\forall F_i \in \Delta^{\Sigma}_{de}(\hat{F})$, $(a,b) \in Att(F_i)$ (i.e. the attack (a,b) belongs to all AFs in the output) meaning that the attack is accepted for a strict majority of agents [10, Definition 38]. When we increase the number of agents that agree with this attack, the majority of agents which support this attack a fortiori*: $\forall G_i \in \Delta^{\Sigma}_{de}(\hat{G}), (a,b) \in Att(G_i)$.*

In what follows, we propose several counter-examples regarding the properties that are not satisfied by Δ^{Σ}_{de}.

Counter-example 1 (σ-SNT, σ-WNT, σ-SD, σ-WD) *We show that even if all the input systems are non-trivial (resp. decisive) we can obtain a trivial (resp. non-decisive) system in the output. We show this by using the* *profile* $\langle F_1, F_2, F_3\rangle$ *depicted in Fig. 8.*

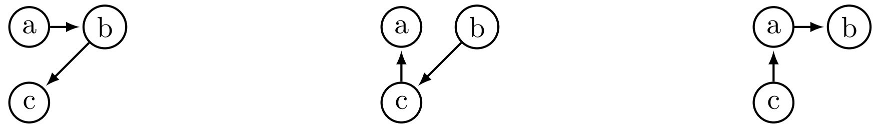

$\mathcal{E}_{pref}(F_1) = \{\{a, c\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{a, b\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{b, c\}\}$

Δ^{Σ}_{de} returns a unique argumentation system:

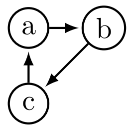

F'

Fig. 8. Δ^{Σ}_{de} falsifies σ-SNT, σ-WNT, σ-SD and σ-WD

Note that each of the three systems in the profile $\langle F_1, F_2, F_3\rangle$ *is non-trivial and decisive. By applying the merging method* Δ^{Σ}_{de}*, we obtain a unique result:* $\Delta^{\Sigma}_{de}(F_1, F_2, F_3) = \{F'\}$ *with* $\sum_{i=1}^{3} de(F', F_i) = 3$ *However, the result of the merging* F' *is a trivial and non-decisive system since* $\mathcal{E}_{pref}(F') = \{\emptyset\}$*. The result is identical for* $\sigma \in \{comp, gr\}$*. For the stable semantics, we obtain* $\mathcal{E}_{sta}(F') = \emptyset$*.*

Counter-example 2 (Unanimity, Majority, Closure) *To show that* Δ^{Σ}_{de} *does not satisfy the properties of Unanimity (*σ*-U,* ca_σ*-U,* sa_σ*-U), Majority (*σ*-MAJ,* ca_σ*-MAJ,* sa_σ*-MAJ) and Closure (CLO,* σ*-C,* ca_σ*-C,* sa_σ*-C), consider the profile* $\langle F_1, F_2, F_3\rangle$ *depicted in Fig. 9.*

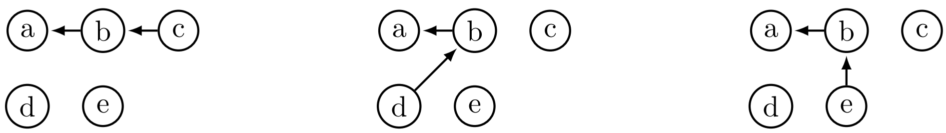

$\mathcal{E}_{pref}(F_1) = \{\{a, c, d, e\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{a, c, d, e\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{a, c, d, e\}\}$

Δ^{Σ}_{de} returns a unique argumentation system:

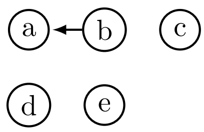

$\mathcal{E}_{pref}(F') = \{\{b, c, d, e\}\}$

Fig. 9. Δ^{Σ}_{de} falsifies the properties of Unanimity, Majority and Closure

By applying Δ^{Σ}_{de}*, we obtain* $\Delta^{\Sigma}_{de}(F_1, F_2, F_3) = \{F'\}$ *with* $\sum_{i=1}^{3} de(F', F_i) = 3$*. However, we have* $\mathcal{E}_{pref}(F') = \{\{b, c, d, e\}\}$ *and hence* $ca_\sigma(F') = sa_\sigma(F') = \{b, c, d, e\}$*.*

This contradicts the properties of Closure since the extension $\{b, c, d, e\}$ does not belong to any AF in the profile (idem for the argument b which is neither skeptically nor credulously accepted by the AFs in the input).
The idea is similar for the properties of Majority since the extension $\{a, c, d, e\}$, even if present in a majoritarian way in the input, does not appear in the extensions of the output framework (idem for the argument a which is skeptically and credulously accepted by a majority of AFs, and not accepted by the result of the merging).
The same reasoning holds for the property of Unanimity.
The result is identical for every semantics $\sigma \in \{comp, pref, sta, gr\}$.

■

Let us check now if there are more properties satisfied when the leximax is used as the aggregation function.

Proposition 4 *Let $\sigma \in \{comp, pref, sta, gr\}$ be a semantics. $\Delta_{de}^{leximax}$ satisfies Anonymity (ANON), the properties of Identity (A-ID, σ-ID, ca_σ-ID, sa_σ-ID), Unanimous attack (A-U), Attack closure (A-C) and Positive responsiveness (PR). The other properties are not satisfied.*

Proof. *For the properties of Anonymity and Identity, the proof is the same as for Δ_{de}^{Σ}.*
Concerning the property of Unanimous attack (A-U), let us suppose that an attack is in all the n argumentation frameworks in the input but is not in the resulting argumentation framework F' which has the minimal score of $\mathcal{L}eximax_{i=1}^{n} de(F', F_i) = (e_1, \ldots, e_n)$. If we add this attack to F', then it obtains a better score (with respect to the leximax) which is $(e_1 - 1, \ldots, e_n - 1)$.
For the property Attack closure (A-C), let us suppose that an attack is in F', which has a minimal score of $\mathcal{L}eximax_{i=1}^{n} de(F', F_i) = (e_1, \ldots, e_n)$, but this attack is in no argumentation framework in the input. If we remove this attack from F', then the new AF obtains a better score $(e_1 - 1, \ldots, e_n - 1)$.
Let us now prove PR. Let $\hat{F}$ and $\hat{G}$ be two profiles such that $\{F_i \in \hat{F} : (a, b) \in Att(F_i)\} \subseteq \{G_i \in \hat{G} : (a, b) \in Att(G_i)\}$, and let for every $F^* \in \gamma(\hat{F})$, $(a, b) \in Att(F^*)$. For every argumentation framework F^* that contains (a, b) its leximax score with respect to $\hat{G}$ is better than or equal to its leximax score with respect to $\hat{F}$. For every argumentation framework F^* that does not contain (a, b) its leximax score with respect to $\hat{G}$ is worse than or equal to its leximax score with respect to $\hat{F}$. Thus, it must be that for every $G^* \in \gamma(\hat{G})$, $(a, b) \in Att(G^*)$.

To see that $\Delta_{de}^{leximax}$ does not satisfy σ-SNT, σ-WNT, σ-SD and σ-WD, consider the example from Fig. 8.

To see that $\Delta_{de}^{leximax}$ does not satisfy σ-U, ca_σ-U, sa_σ-U, σ-MAJ, ca_σ-MAJ, sa_σ-MAJ, CLO, σ-C, ca_σ-C and sa_σ-C, consider the example from Fig. 9.

Counter-example 3 (A-MAJ) *Fig. 10 contains a counter-example showing that Majority attack is not satisfied by* $\Delta_{de}^{leximax}$.

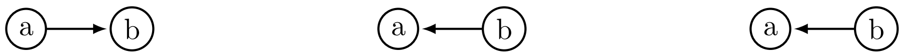

$\mathcal{E}_{pref}(F_1) = \{\{a\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{b\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{b\}\}$

$\Delta_{de}^{leximax}$ returns the following argumentation frameworks:

F'_1 $\quad$ F'_2

Fig. 10. $\Delta_{de}^{leximax}$ falsifies Majority attack

The property A-MAJ is not satisfied because the attack (b, a) *is not in* F'_2.

■

Let us now turn to the properties satisfied by the qualified voting method [27]. Note that we suppose that qualified voting satisfies a given postulate if and only if it satisfies this postulate for all $|U| \geq 2$ from definition 11.

Proposition 5 *QV satisfies gr-weak non triviality (gr-WNT), gr-weak decisiveness (gr-WD), Attack closure (A-C) and Positive responsiveness (PR) for every* $|U| \geq 2$. *The other properties are not satisfied.*

Proof. *The proof that QV satisfies PR can be found in [27].*
Regarding A-C, the proof comes from the definition of QV. Indeed, an attack appears in the result of the merging if and only if there are more agents that agree with that attack than those who disagree (i.e. those who consider that there is no attack or that there is an attack in the inverse direction).
For the proof of gr-WD, we recall that a non-empty grounded extension exists if there exists at least one non-attacked argument in the AF. Let $\hat{F} = \langle F_1, \ldots, F_n \rangle$ *be a profile and* $U \subset \{1, \ldots, n\}$. *Recall that* $\forall i \; F_i \in \mathbb{AF}_{D_{gr}}$ *so there exists at least one non-attacked argument in each* F_i. *Let us prove that each non-attacked argument (proposed by each agent in U) stays non-attacked after the merging. Hence, the result, noted* $\gamma(\hat{F})$, *will be decisive too.*
Let an argument $x \in Arg(F_1)$ *such that* x *is non-attacked in* F_1 *and* $1 \in U$. *Suppose that there exists an argument* $y \in Arg(\gamma(\hat{F}))$ *such that* $(y, x) \in Att(\gamma(\hat{F}))$. *That means that all the agents in U should agree with this attack (see the second condition of the definition of QV). However, it is not the case of* F_1 *because* x *is non-attacked. So* x *stays non-attacked in* $\gamma(\hat{F})$ *implying the*

existence of a non-empty grounded extension which is unique by definition. The proof is similar for gr-WNT because of the uniqueness of the grounded extension.

QV does not satisfy the property ANON because of the set of agents U which gives more importance to some argumentation frameworks. In addition, the permutation does not garantee that the set of agents U is correctly rearranged too which can lead to a different result.

Counter-example 4 (gr-SNT, gr-SD) *The example from Fig. 11 shows that, with $U = \{1, 2\}$, gr-SNT and gr-SD are not satisfied by QV because F' is trivial and non-decisive.*

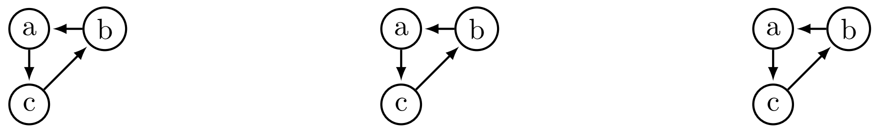

$\mathcal{E}_{gr}(F_1) = \{\emptyset\}$ $\quad$ $\mathcal{E}_{gr}(F_2) = \{\emptyset\}$ $\quad$ $\mathcal{E}_{gr}(F_3) = \{\emptyset\}$

The result of the merging is a trivial and non-decisive unique AF:

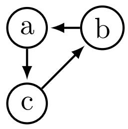

$\mathcal{E}_{gr}(F') = \{\emptyset\}$

Fig. 11. QV falsifies gr-SNT and gr-SD

Counter-example 5 (σ-SNT, σ-WNT, σ-SD, σ-WD) *The example from Fig. 12 shows that, with $U = \{1, 2\}$, QV does not satisfy σ-SNT, σ-WNT, σ-SD and σ-WD for $\sigma \in \{pref, sta, comp\}$ because F' is trivial and non-decisive.*

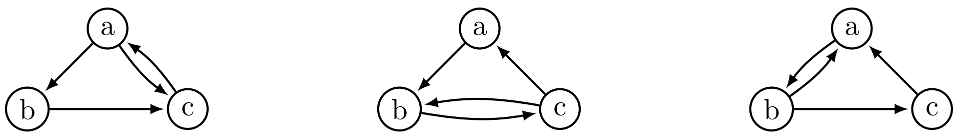

$\mathcal{E}_{pref}(F_1) = \{\{a\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{c\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{b\}\}$

Applying QV gives a trivial and non-decisive unique AF:

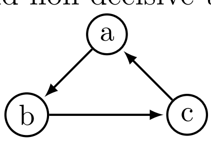

$\mathcal{E}_{pref}(F') = \{\emptyset\}$

Fig. 12. QV falsifies the properties of Non-Triviality and Decisiveness

It is easy to see that the same counter-example can be used for stable and complete semantics.

Counter-example 6 (Unanimity, Majority, Closure) *To show that the properties σ-U, ca_σ-U, sa_σ-U, σ-MAJ, ca_σ-MAJ, sa_σ-MAJ, σ-C, ca_σ-C and sa_σ-C are violated, we can use the counter-example from Fig. 9 for any $U \subseteq \{1, 2, 3\}$ and $|U| \geq 2$. Indeed, the only attack resulting of the aggregation of the three AFs is the attack from b to a. Thus, we obtain, for the four classical semantics, one extension which is different from the extension obtained in each AF in the input.*

Counter-example 7 (Identity, A-ID, A-MAJ, A-U, CLO) *Consider the example from Fig. 13, with $U = \{1, 2\}$, to show that σ-ID, ca_σ-ID, sa_σ-ID, A-ID, A-MAJ, A-U and CLO are not satisfied by QV.*

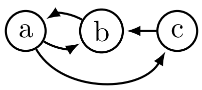

$\mathcal{E}_{pref}(F_1) = \{\{a\}\}$

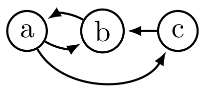

$\mathcal{E}_{pref}(F_2) = \{\{a\}\}$

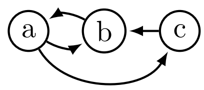

$\mathcal{E}_{pref}(F_3) = \{\{a\}\}$

By applying QV, we obtain a unique result:

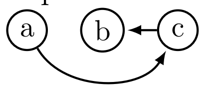

$\mathcal{E}_{pref}(F') = \{\{a, b\}\}$

Fig. 13. QV falsifies the Identity properties, A-MAJ, A-U and CLO

The same counter-example can be used for $\sigma \in \{comp, pref, sta, gr\}$.

■

7 Using WAFs for Aggregating AFs

Let us now propose new aggregation methods based on WAFs. When WAFs were introduced [20,13,14] one of the possible interpretations of the weights on an attack was that it could represent the number of agents in a group that agree with this attack. So we endorse this interpretation and study how we can define operators that aggregate a set of AFs using techniques developed for WAFs.

7.1 $\mathbf{FUS}_{All}$

The first method, noted FUS_{All}, consists in simply building a WAF where the weights represent the number of agents that agree with (i.e. that have) a given attack. Once built, we use one of four *best* methods (see Definition 9) in order to obtain a set of extensions representing the result of the aggregation of the profile.

Definition 12 *Let $\hat{F} = (F_1, \ldots, F_n)$ be a profile*

$$\mathrm{FUS}_{All}^{best_i^{\sigma,\oplus}}(\hat{F}) = best_i^{\sigma,\oplus}(waf(\hat{F}))$$

where $waf(\hat{F}) = \langle A, R, w\rangle$, with:

- $A = \mathcal{X}$,
- $R = \bigcup_{i=1}^{n} Att(F_i)$,
- *and* $w(a, b) = |\{F_i \in \hat{F} : (a, b) \in Att(F_i)\}|$.

Note that the construction of $waf(\hat{F})$ is exactly the one proposed by Cayrol and Lagasquie-Schiex [9] up to a normalization of the weights, but nothing is said about what to do with the obtained WAF. We propose to use the *best* methods in order to find extensions as an output.

Example 3 (cont.) *From the profile illustrated in Fig. 5, let us begin to build the corresponding WAF (see Fig. 14).*

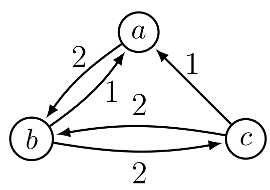

Fig. 14. WAF obtained with FUS_{All}

In using the preferred semantics, there are two distinct extensions $\mathcal{E}_{pref}(\overline{\mathrm{WF}}) = \{\{b\}, \{c\}\}$ which cannot be distinguished with the best methods because the weight on the attack from b to c and the weight on the attack from c to b are identical. So, $\mathrm{FUS}_{All}^{best_i^{pref,\oplus}}(\langle F_1, F_2, F_3\rangle) = \{\{b\}, \{c\}\}$.
Conversely, with the grounded semantics, there exists no non-empty extension from $\overline{\mathrm{WAF}}$: $\mathcal{E}_{gr}(\overline{\mathrm{WF}}) = \{\emptyset\}$. With no extension to compare, the result of aggregation stays the same for all best method: $\mathrm{FUS}_{All}^{best_i^{gr,\oplus}}(\langle F_1, F_2, F_3\rangle) = \{\emptyset\}$.

Let us now check which properties are satisfied by FUS_{All}. Note that our operators produce as a result a set of extensions. Hence, some of the properties, namely Unanimous attack (A-U), Majority attack (A-MAJ), Closure (CLO), Attack closure (A-C), Identity attack (A-ID) and Positive Responsiveness (PR), dealing with the attacks relation, are not applicable here. We recall also that in this paper we focus on the main semantics defined by Dung : $\sigma \in \{comp, pref, sta, gr\}$. Finally, concerning the best extensions (see Definition 9), we choose to study the four best rules with the sum and the max as aggregation function ($\oplus \in \{\Sigma, max\}$).

Proposition 6 *Let $\sigma \in \{comp, pref, sta, gr\}$ be a semantics.* FUS_{All} *satisfies Anonymity (ANON) and properties gr-Identity (gr-ID), ca_{gr}-Identity (ca_{gr}-ID) and sa_{gr}-Identity (sa_{gr}-ID) for each best rule and for each aggregation function $\oplus \in \{\Sigma, max\}$. The other properties are not satisfied.*

Proof. *The claim for Anonymity is obvious since* FUS_{All} *does not use a specific order for the aggregation. Concerning the properties of Identity, the three properties based on extensions are satisfied because of the uniqueness of the grounded extension.*

Counter-example 8 (σ-SNT, σ-WNT, σ-SD, σ-WD) *Consider the example from Fig. 15 where F_1, F_2 and F_3 are non-trivial and decisive.*

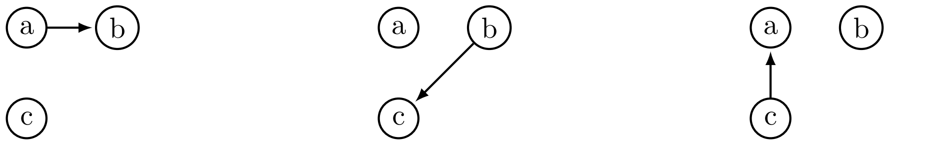

$\mathcal{E}_{pref}(F_1) = \{\{a, c\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{a, b\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{b, c\}\}$

We obtain the following trivial and non-decisive WAF:

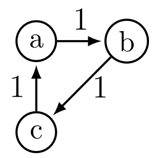

$$best_1^{pref,\Sigma}(\mathrm{WF}) = \{\emptyset\}$$

Fig. 15. *FUS_{All}* falsifies σ-SNT, σ-WNT, σ-SD and σ-WD

The same example can be used with other best *methods and for every semantics $\sigma \in \{comp, pref, sta, gr\}$.*

Counter-example 9 (gr-U, ca$_{gr}$-U, sa$_{gr}$-U) *To show that the properties of Unanimity (gr-U, sa_{gr}-U, ca_{gr}-U) are not satisfied under grounded semantics, consider the example from Fig. 16.*

$$\mathcal{E}_{gr}(F_1) \cap \mathcal{E}_{gr}(F_2) = \{a, c, e\} \nsubseteq \{\{a\}\} = best_1^{gr,\Sigma}(\mathrm{WF})$$

Consequently, e and c, which are skeptically and credulously satisfied in all AFs in the input, are not in the result of the aggregation. We can note that this counter-example gives a similar result if another best *method is used because of the uniqueness of the result.*

Counter-example 10 (σ-U, ca_σ-U, σ-ID, ca_σ-ID, sa_σ-ID) *We show that σ-U, ca_σ-U, σ-ID, ca_σ-ID and sa_σ-ID are falsified for $\sigma \in \{sta, pref, comp\}$.*

Suppose three identical argumentation systems F_1, F_2, F_3 (with the same ar-

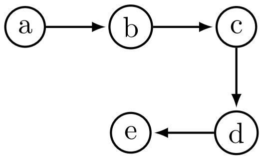

$\mathcal{E}_{gr}(F_1) = \{a, c, e\}$

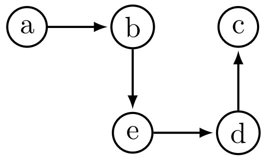

$\mathcal{E}_{gr}(F_2) = \{a, c, e\}$

We obtain the following WAF:

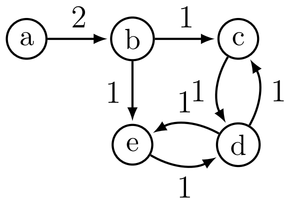

$best_1^{gr,\Sigma}(\mathrm{WF}) = \{a\}$

Fig. 16. *FUS*$_{All}$ falsifies gr-U, ca_{gr}-U and sa_{gr}-U

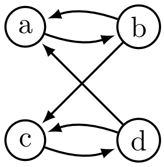

$\mathcal{E}_{pref}(F_{1,2,3}) = \{\{a, c\}, \{b, d\}\}$

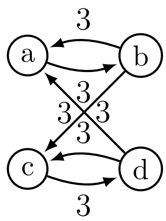

$best_1^{pref,\Sigma}(\mathrm{WF}) = \{\{b, d\}\}$

Fig. 17. FUS_{All} falsifies σ-U, ca_σ-U, σ-ID, ca_σ-ID and sa_σ-ID

guments and the same attack relation), depicted on the left side of Fig. 17. Note that $\mathcal{E}_{pref}(F_{1,2,3}) = \{\{a, c\}, \{b, d\}\}$.
The corresponding WAF is on the right side of Fig. 17. We obtain $\mathcal{E}_{pref}(\overline{\mathrm{WF}}) = \{\{a, c\}, \{b, d\}\}$.
Let us compare the two extensions using $best_1^\Sigma$.

$$\left.\begin{aligned} S_\Sigma(\{a,c\} \to \{b,d\}) &= w(a,b) + w(a,d) + \ldots = 6 \\ S_\Sigma(\{b,d\} \to \{a,c\}) &= w(b,a) + w(b,c) + \ldots = 12 \end{aligned}\right\} \{a,c\} <_\Sigma \{b,d\}$$

We have $best_1^{pref,\Sigma}(\mathrm{WF}) = \{\{b, d\}\}$.

- σ*-U,* σ*-ID :* $\{a, c\} \not\subseteq \{\{b, d\}\} = best_1^{pref,\Sigma}(\mathrm{WF})$
- ca_σ*-U,* ca_σ*-ID :* $a \notin \{b, d\} = ca_{pref}(best_1^{pref,\Sigma}(\mathrm{WF}))$ *(idem for c)*
- sa_σ*-ID :* $sa_{pref}(F_{1,2,3}) = \emptyset \neq \{b, d\} = sa_{pref}(best_1^{pref,\Sigma}(\mathrm{WF}))$

Using leximax instead of sum and/or another best method yields the same result. The same counter-example can be used for stable and complete semantics.

Counter-example 11 (sa_σ-U) *Consider the systems from Fig. 16. F_1 and F_2 have the same extensions $\mathcal{E}_{pref}(F_1) = \mathcal{E}_{pref}(F_2) = \{\{a, c, e\}\}$, and consequently $sa_{pref}(F_1) = sa_{pref}(F_2) = \{a, c, e\}$.*
The extensions of the corresponding $\overline{\mathrm{WF}}$ are: $\mathcal{E}_{pref}(\overline{\mathrm{WF}}) = \{\{a, d\}, \{a, c, e\}\}$. We have $best_1^{pref,\Sigma}(\mathrm{WF}) = \{\{a, d\}, \{a, c, e\}\}$, so $sa_{pref}(best_1^{pref,\Sigma}(\mathrm{WF})) = \{a\}$. So we obtain:

$$sa_{pref}(F_1) \cap sa_{pref}(F_2) = \{a, c, e\} \nsubseteq \{a\} = sa_{pref}(best_1^{pref,\Sigma}(\mathrm{WF}))$$

Applying stable or complete semantics together with another best *method (or max instead of sum) gives a similar result.*

Counter-example 12 (σ-MAJ, ca_σ-MAJ, sa_σ-MAJ) *Consider the example depicted in Fig. 18.*

$\mathcal{E}_{pref}(F_1) = \{\{a, b\}\}$ $\quad$ $\mathcal{E}_{pref}(F_2) = \{\{a, b\}\}$ $\quad$ $\mathcal{E}_{pref}(F_3) = \{\{a\}\}$

We obtain the following trivial WAF :

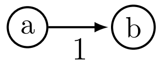

$best_1^{pref,\Sigma}(\mathrm{WF}) = \{\{a\}\}$

Fig. 18. FUS_{All} falsifies σ-MAJ, ca_σ-MAJ and sa_σ-MAJ

- *σ-MAJ is not satisfied since the extension $\{a, b\}$, present in two input systems, does not appear as an extension of the resulting system: $\{a, b\} \notin \{\{a\}\} = best_1^{pref,\Sigma}(\mathrm{WF})$*
- *ca_σ-MAJ is not satisfied because b is credulously accepted in a strict majority of AF in the input, whereas it does not appear in the set of credulous arguments of the resulting system: $b \notin \{a\} = ca_{pref}(best_1^{pref\Sigma}(\mathrm{WF}))$*
- *We use the same reasoning for sa_σ-MAJ: $b \notin \{a\} = sa_{pref}(best_1^{pref,\Sigma}(\mathrm{WF}))$*

The same example can be used for $\sigma \in \{comp, pref, sta, gr\}$ and other best *methods.*

Counter-example 13 (σ-C, ca_σ-C, sa_σ-C) *Consider the example from Fig. 19.*

- *σ-closure : $\{a, c\} \notin \{\{a, d\}, \{a, b\}\} = \mathcal{E}_{gr}(F_1) \cup \mathcal{E}_{gr}(F_2)$*
- ca_σ-closure, sa_σ-closure *: c is credulously (resp. skeptically) accepted in the result of the aggregation. However, it is not credulously (resp. skeptically) accepted in any of the input systems : $c \notin \{a, b, d\} = ca_{gr}(F_1) \cup ca_{gr}(F_2) = sa_{gr}(F_1) \cup sa_{gr}(F_2)$.*

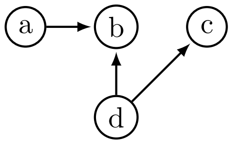

$\mathcal{E}_{gr}(F_1) = \{a, d\}$

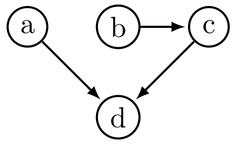

$\mathcal{E}_{gr}(F_2) = \{a, b\}$

We obtain the following WAF :

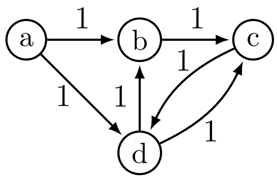

$best_1^{gr,\Sigma}(\text{WF}) = \{\{a, c\}\}$

Fig. 19. FUS_{All} falsifies gr-C, ca_{gr}-C and sa_{gr}-C

The same example can be used for $\sigma \in \{comp, pref, sta, gr\}$*. As the result contains exactly one extension, this result is the same for any* best *method used.*

■

In particular, one of the properties that is not satisfied is non-triviality, meaning that these operators do not ensure the existence of at least one result in the output, which can be considered as an important drawback.

7.2 FUS$_{AllNT}$

A solution to satisfy non-triviality, i.e. to ensure that the set of extensions of a $\overline{\text{WF}}$ is always non-empty, is to use the relaxed extensions techniques (see Section 3.1).

Definition 13 *Let* $\hat{F} = (F_1, \dots, F_n)$ *be a profile*

$$\text{FUS}_{AllNT}^{\sigma, best_i^{\oplus}, \otimes}(\hat{F}) = best_i^{\oplus}(\mathcal{E}_\sigma^{\otimes}(\overline{waf(\hat{F})}, waf(\hat{F})))$$

where $waf(\hat{F}) = \langle A, R, w \rangle$*, with:*

- $A = \mathcal{X}$,
- $R = \bigcup_{i=1}^{n} Att(F_i)$,
- *and* $w(a, b) = |\{F_i \in \hat{F} : (a, b) \in Att(F_i)\}|$.

Concerning the aggregation function used for relaxed extensions, we only focus on the sum [4] ($\otimes = \Sigma$).

[4] That is the original definition by Dunne et al. [21].

Example 3 (cont.) *From the profile illustrated in Fig. 5, we obtain the same WAF that* FUS_{All} *(see Fig. 14). Concerning the preferred semantics, the result is exactly the same as* FUS_{All} *because the weighted argumentation framework obtained is non-trivial. So,* $\mathrm{FUS}_{AllNT}^{best_i^{pref,\oplus},\otimes}(\langle F_1, F_2, F_3\rangle) = \{\{b\}, \{c\}\}$.
With no non-empty grounded extension, we have to remove some attacks to find at least one non-empty extension. The smallest β *which returns a non-empty extension with the sum as aggregation function is 2 because we obtain two extensions:* $\mathcal{E}_{gr}^{\Sigma,2}(\mathrm{WF}) = \{\{c\}, \{a, c\}\}$. *Indeed, either the attack from b to c with a weight of 2 is removed to obtain the extension* $\{c\}$, *or we remove the two attacks with a weight of 1 to obtain the extension* $\{a, c\}$. *However, it is possible to make a distinction between the two extensions because c attacks a,* i.e. $S_\Sigma(\{a, c\} \rightarrow \{c\}) = 0$ *while* $S_\Sigma(\{c\} \rightarrow \{a, c\}) = w(c, a) = 1$. *Consequently, with the best methods, we have* $\mathrm{FUS}_{AllNT}^{best_i^{gr,\oplus},\Sigma}(\langle F_1, F_2, F_3\rangle) = \{\{c\}\}$.

Proposition 7 *Let* $\sigma \in \{comp, pref, sta, gr\}$ *be a semantics. Let* $\otimes = \Sigma$. FUS_{AllNT} *satisfies Anonymity (ANON),* σ*-strong non-triviality (*σ*-SNT),* σ*-weak non-triviality (*σ*-WNT) and properties gr-Identity (gr-ID),* ca_{gr}*-Identity (*ca_{gr}*-ID) and* sa_{gr}*-Identity (*sa_{gr}*-ID) for each best rule and each aggregation function* $\oplus \in \{\Sigma, max\}$. *The other properties are not satisfied.*

Proof. *The proof for Anonymity follows directly from the definition. For the properties of non-triviality, the proof is also trivial, since the relaxing mechanisms guarantees the existence of at least one extension. Let us show that the gr-Identity,* ca_{gr}*-Identity and* sa_{gr}*-Identity are satisfied. If the profile contains n identical argumentation frameworks, the merged framework has the same grounded extension and the best method is not used (since the grounded extension is unique). Furthermore, since the definition of* σ*-Identity,* ca_σ*-Identity and* sa_σ*-Identity require the input frameworks to be non-trivial, their grounded extensions are not empty. Hence, the grounded extension of the output framework is not empty too.*

Regarding all properties that are not satisfied by FUS_{AllNT}, *some counter-example used for* FUS_{All} *can also be used for* FUS_{AllNT}. *It is the case of the properties* σ*-U,* ca_σ*-U,* sa_σ*-U,* σ*-ID,* ca_σ*-ID,* sa_σ*-ID,* σ*-MAJ,* ca_σ*-MAJ,* sa_σ*-MAJ,* σ*-C,* ca_σ*-C and* sa_σ*-C. Indeed, none of those examples had a trivial result, so using* FUS_{AllNT} *does not change anything.*

Counter-example 14 (*σ*-SD,* σ*-WD) *The same counter-example as for* FUS_{All} *(see Fig. 15) can be used. Indeed, in using the relaxed extensions, we obtain the following result which is not decisive too :*

$$best_1^{\Sigma}(\mathcal{E}_{pref}^{\Sigma}(\overline{\mathrm{WF}}, \mathrm{WF})) = \{\{a, b\}, \{a, c\}, \{b, c\}\}$$

The same example can be used with other best *methods and for every seman-*

tics $\sigma \in \{comp, pref, sta, gr\}$.

■

This operator is useful if we want to take into account all the attacks given by the agents. However, the result does not ensure the representation of the opinion of the majority. For instance, suppose that we have nine AFs with $A = \{a, b\}$ and $R = \{\}$ and one AF with the same set of arguments but a attacks b ($R = \{(a, b)\}$). If we merge these ten AFs by using FUS_{All} and FUS_{AllNT}, then the attack (a, b), only given by one agent, is present in the resulting system. This is clearly against the opinion of the majority.

7.3 $\mathbf{FUS}_{MajNT}$

A more natural way of constructing the WAF corresponding to the set of AFs should take into account the notion of *majority* during the construction of the WAF. This means that, instead of representing all the attacks of the profile, we only select the attacks accepted by a strict majority of agents.

Definition 14 *Let* $\hat{F} = (F_1, \ldots, F_n)$ *be a profile.*

$$\mathrm{FUS}_{MajNT}^{\sigma, best_i^{\oplus}, \otimes}(\hat{F}) = best_i^{\oplus}(\mathcal{E}_{\sigma}^{\otimes}(\overline{mwf(\hat{F})}, mwf(\hat{F})))$$

where $mwf(\hat{F}) = \langle A, R, w \rangle$, *with:*

- $A = \mathcal{X}$,
- $R = \{(a, b) : |\{F_i : (a, b) \in Att(F_i)\}| > \frac{n}{2}\}$,
- *and* $w(a, b) = |\{F_i \in \hat{F} : (a, b) \in Att(F_i)\}|$ *if* $(a, b) \in R$, *and* $= 0$ *otherwise.*

Example 3 (cont.) *From the profile illustrated in Fig. 5, we obtain the WAF illustrated in Fig. 20 that contains the three attacks that appear in a strict majority of AFs in the input.*

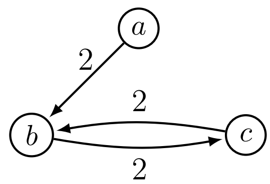

Fig. 20. WAF obtained with FUS_{MajNT}

The grounded and the preferred extensions are the same here : $\mathcal{E}_{pref}(\overline{\mathrm{WF}}) = \mathcal{E}_{gr}(\overline{\mathrm{WF}}) = \{\{a, c\}\}$. *This extension is non-trivial (so we do not need to find* *the relaxed extensions) and is the only one (so this is obviously the best extension).* $\mathrm{FUS}_{MajNT}^{best_i^{pref,\oplus},\otimes}(\langle F_1, F_2, F_3\rangle) = \mathrm{FUS}_{MajNT}^{best_i^{gr,\oplus},\otimes}(\langle F_1, F_2, F_3\rangle) = \{\{a, c\}\}$.

Let us now check what properties are satisfied by FUS_{MajNT}.

Proposition 8 *Let* $\sigma \in \{comp, pref, sta, gr\}$ *be a semantics. Let* $\otimes = \Sigma$. FUS_{MajNT} *satisfies Anonymity (ANON),* σ*-strong non-triviality (*σ*-SNT),*σ*-weak non-triviality (*σ*-WNT) and properties gr-Identity (gr-ID),* ca_{gr}*-Identity (*ca_{gr}*-ID) and* sa_{gr}*-Identity (*sa_{gr}*-ID) for each best rule and each aggregation function* $\oplus \in \{\Sigma, max\}$*. The other properties are not satisfied.*

Proof. *The proofs are similar to the ones for* FUS_{AllNT}.

Counter-example 15 (*σ*-SD,* σ*-WD) *Consider the example from Fig. 21.*

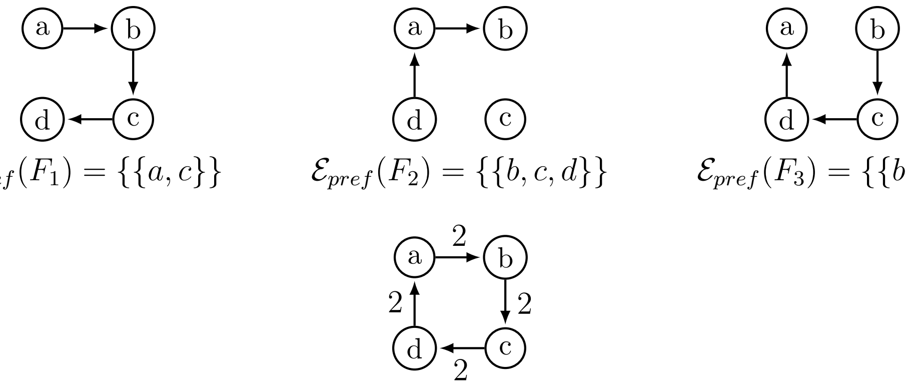

$\mathcal{E}_{pref}(F_1) = \{\{a, c\}\}$ $\mathcal{E}_{pref}(F_2) = \{\{b, c, d\}\}$ $\mathcal{E}_{pref}(F_3) = \{\{b, d\}\}$

$best_1^{\Sigma}(\mathcal{E}_{pref}^{\Sigma}(\mathrm{WF}), \mathrm{WF}) = \{\{a, c\}, \{b, d\}\}$

Fig. 21. FUS_{MajNT} falsifies σ-SD and σ-WD

Note that each of F_1, F_2, F_3 *is decisive. However, the result of the merging is not decisive. Both extensions are equally preferred by each of* best *methods. The same example can be used for other semantics. For the grounded semantics, we obtain this result thanks to the relaxed extensions.*

Counter-example 16 (Unanimity, Majority, Closure) *Let us show that* FUS_{MajNT} *violates* σ*-U,* ca_σ*-U,* sa_σ*-U,* σ*-MAJ,* ca_σ*-MAJ,* sa_σ*-MAJ,* σ*-C,* ca_σ*-C and* sa_σ*-C. Let* F_1*,* F_2 *and* F_3 *be three AFs represented in Fig. 22. Each of them has a unique grounded extension:* $\{a, c, d, e\}$.

The corresponding WAF is depicted in the same figure, below the three systems. Applying FUS_{MajNT} *yields a unique extension* $\{b, c, d, e\}$ *(see Section 8 for more details). The same example can be used with other* best *methods and for every semantics* $\sigma \in \{comp, pref, sta, gr\}$.

■

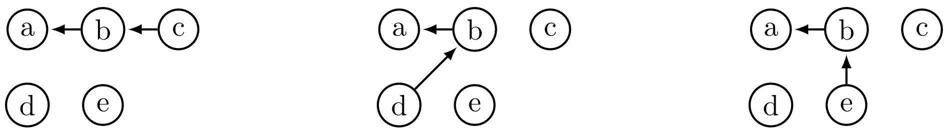

$\mathcal{E}_{gr}(F_1) = \{\{a, c, d, e\}\}$ $\quad$ $\mathcal{E}_{gr}(F_2) = \{\{a, c, d, e\}\}$ $\quad$ $\mathcal{E}_{gr}(F_3) = \{\{a, c, d, e\}\}$

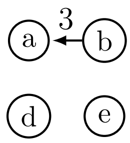

$$best_1^{\Sigma}(\mathcal{E}_{gr}^{\Sigma}(\mathrm{WF}), \mathrm{WF}) = \{\{b, c, d, e\}\}$$

Fig. 22. FUS_{MajNT} falsifies σ-U, ca_σ-U, sa_σ-U, σ-MAJ, ca_σ-MAJ, sa_σ-MAJ, σ-C, ca_σ-C, sa_σ-C

Note the surprising fact that the properties of *majority* are not satisfied by FUS_{MajNT}. We will explain the reasons in the next section.

8 Discussion

Let us sum up our results in Table 2 and discuss their impact.
It is clear that there are few properties satisfied by the existing aggregation operators. There are two possible (non-exclusive) explanations: either the existing operators are not good enough, or the "rationality" properties are too demanding. Our point of view is that both are true to some extent. Indeed, more work is needed both in defining a set of rationality properties that capture more adequately the desirable behaviour of an aggregation operator, and on defining aggregation methods themselves.

Let us first argue that some of the properties are too strong (this means that not satisfying them is not a disqualifying feature for an aggregation operator). Let us recall the example from Fig. 9.

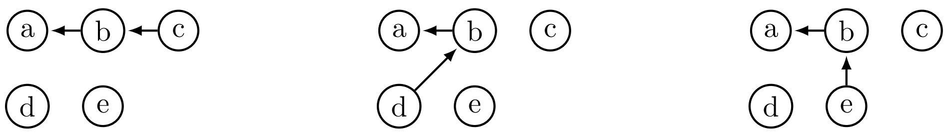

F_1 $\quad$ F_2 $\quad$ F_3

So, if one focuses on the attack relation, the intuitive output is the following AF:

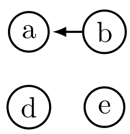

Properties	$\triangle_{de}^{\Sigma}$	$\triangle_{de}^{leximax}$	QV	FUS_{All}	FUS_{AllNT}	FUS_{MajNT}
ANON	✓	✓	×	✓	✓	✓
σ-SNT	×	×	×	×	✓	✓
σ-WNT	×	×	$\checkmark^{gr}$	×	✓	✓
σ-SD	×	×	×	×	×	×
σ-WD	×	×	$\checkmark^{gr}$	×	×	×
A-U	✓	✓	×	-	-	-
σ-U	×	×	×	×	×	×
ca_{σ}-U	×	×	×	×	×	×
sa_{σ}-U	×	×	×	×	×	×
A-MAJ	✓	×	×	-	-	-
σ-MAJ	×	×	×	×	×	×
ca_{σ}-MAJ	×	×	×	×	×	×
sa_{σ}-MAJ	×	×	×	×	×	×
CLO	×	×	×	-	-	-
A-C	✓	✓	✓	-	-	-
σ-C	×	×	×	×	×	×
ca_{σ}-C	×	×	×	×	×	×
sa_{σ}-C	×	×	×	×	×	×
A-ID	✓	✓	×	-	-	-
σ-ID	✓	✓	×	$\checkmark^{gr}$	$\checkmark^{gr}$	$\checkmark^{gr}$
ca_{σ}-ID	✓	✓	×	$\checkmark^{gr}$	$\checkmark^{gr}$	$\checkmark^{gr}$
sa_{σ}-ID	✓	✓	×	$\checkmark^{gr}$	$\checkmark^{gr}$	$\checkmark^{gr}$
PR	✓	✓	✓	-	-	-

Table 2

Properties × Aggregation operators. A cross × means that the property is not satisfied, symbol ✓ means that the property is satisfied, $\checkmark^{\sigma}$ means that the property is satisfied for the semantics σ, and symbol $-$ means that the property can not be applied to the operator (because the output of the operator is not compatible with the constraint given by the rule)

This proposed result can seem, at first sight, illogical if we focus on extensions (and consequently the accepted arguments) since each F_i has the same extension $\{a, c, d, e\}$, so we could expect this extension to be the outcome. However, if we look more closely at this example, we can see that this same extension is obtained for very different reasons. Each agent has a reason (argument) to reject b, but this attack is challenged by all the other agents, and thus could be interpreted as an error of this agent. So it is quite natural to refuse all attacks on b for the outcome aggregation framework, which means that $\{a, c, d, e\}$ should be rejected as the extension of the outcome (and that $\{b, d, c, e\}$ is much more natural). In fact, this AF is obtained by using the majority method (majority vote on the attack relation): all the F_i agree on the attack (b, a), whereas all other attacks have a maximum of one F_i supporting it. This is also the result obtained when FUS_{MajNT} is used. However, that goes against the properties of unanimity, majority and closure related to extensions, credulous inference and skeptical inference [5]. So if one wants to

[5] σ-MAJ, ca_{σ}-MAJ, sa_{σ}-MAJ, σ-U, ca_{σ}-U, sa_{σ}-U, CLO, σ-C, ca_{σ}-C, and sa_{σ}-C

obtain the expected outcome of this example, only the properties about the attack relation seem to not be problematic.
We want to insist on the fact that we give here a simple example with only three argumentation frameworks, but this example can be generalized with 100 agents (and 102 arguments), such that (b, a) is supported by all, and each agent i supports only an additional attack between argument a_i and b (and he is the only one to support it). In this case the quasi-unanimity situation (all agents except one are against the other attacks) is much more striking.
The properties of decisiveness seem also much too strong requirements for most semantics that accept several extensions (and are trivial for the ones that accept at most one extension), so we propose to remove them from necessary properties also.

Basically our opinion is that the proposed properties were more or less direct translations of properties coming from social choice theory. This was certainly an important first step. However, argumentation frameworks have more structure than sets of candidates in voting problems, so the specificities of this structure of AFs have to be taken into account. We argue that these structural specificities invalidate some of the properties from social choice theory as being required for aggregation of AFs. This does not mean that they are not of interest, since they can be used to characterize some aggregation methods (there should be some methods that satisfy them), but they can not be considered as absolutely necessary requirements. An alternative is proposed by Delobelle et al. [17] where the postulates from propositional merging are restated in the domain of abstract argumentation and a generic representation theorem is derived for extension-based argumentation merging. These postulates seem more adequate for argumentation than properties coming from social choice theory, since they deal with structured pieces of information.

The shaded rows in Table 2 contain, in our opinion, the most desirable properties if one concentrates on the attack relation. Indeed, ANON seems to be a basic requirement to avoid the problem of giving more importance to some agents. The properties which focus on the attack relation in the result of the aggregation (*i.e.* A-U, A-MAJ, A-ID and PR) should obviously be satisfied because when at least a strict majority of agents agree with an attack, the result should have this attack too in order to correctly reflect the group's point of view. Concerning A-C, it seems difficult to justify the fact that an attack appears in the resulting AF while it does not appear in any AF in the input. We consider σ-WNT to be a mandatory requirement - as it guarantees the existence of a solution of the aggregation when the AFs at the input are non-trivial.

However, one can see that there is no existing aggregation method that fully satisfies all of these properties. This means that there is still work needed to define good aggregation operators.

The previous example illustrates that there seems to be some incompatibilities between the rationality properties for aggregation of argumentation frameworks that deal with extensions and the ones that deal with attacks. Both approaches seem sensible, so there should be two different sets of postulates, depending on the chosen priority, i.e. one that focuses on extensions and one that concentrates on the attack relation. To go further, two families of aggregation operators of argumentation frameworks seem to appear: one focuses on the attack relations like all the operators defined in this paper while the other focuses on the extensions (see operators defined in [17]).

9 Conclusion

In this paper we put together the works from the literature on aggregation methods for Dung's abstract argumentation frameworks. We focus on the methods that take as input a profile of abstract argumentation frameworks, and give as a result an argumentation framework, a set of argumentation frameworks, or a set of extensions. We also investigate the use of WAFs in order to aggregate profiles of AFs, and we end up with three possible definitions, FUS_{MajNT} being certainly the most convincing. We show that few of the proposed properties are satisfied by existing aggregation operators. The explanation seems to incriminate both suspects: the properties and the methods. On one hand, some of the properties seem to be too demanding in the general case. On the other hand, the existing operators do not satisfy even the most desirable properties.
Our results seem to suggest that a lot of work is still needed on the two fronts. A more careful study of the rationality properties for aggregation methods for abstract argumentation is required and there is clearly room for definition of other (possibly better) aggregation methods.

10 Acknowledgement

We thank the reviewers for their useful comments on the previous version of the paper. This work benefited from the support of the project AMANDE ANR-13-BS02-0004 of the French National Research Agency (ANR).

References

[1] L. Amgoud and S. Vesic. On revising argumentation-based decision systems. In *Proc. of the 10th European Conference on Symbolic and Quantitative*

Approaches to Reasoning with Uncertainty (ECSQARU'09), pages 71–82. Springer, 2009.

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

[3] R. Baumann. Normal and strong expansion equivalence for argumentation frameworks. *Artificial Intelligence*, 193:18–44, 2012.

[4] T. J. M. Bench-Capon and P. E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[5] G. Boella, S. Kaci, and L. van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Proc. of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09)*, pages 107–118. Springer, 2009.

[6] E. Bonzon and N. Maudet. On the outcomes of multiparty persuasion. In *Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pages 47–54, 2011.

[7] M. Caminada and G. Pigozzi. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems*, 22(1):64–102, 2011.

[8] C. Cayrol, F. D. de Saint-Cyr, and M. Lagasquie-Schiex. Revision of an argumentation system. In *Proc. of the 11th International Conference on the Principles of Knowledge Representation and Reasoning (KR'08)*, pages 124–134, 2008.

[9] C. Cayrol and M. Lagasquie-Schiex. Weighted argumentation systems: A tool for merging argumentation systems. In *IEEE 23rd International Conference on Tools with Artificial Intelligence (ICTAI'11)*, pages 629–632, 2011.

[10] S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasquie-Schiex, and P. Marquis. On the merging of Dung's argumentation systems. *Artificial Intelligence*, 171(10-15):730–753, 2007.

[11] S. Coste-Marquis, S. Konieczny, J.-G. Mailly, and P. Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In *Proc. of the 14th International Conference on the Principles of Knowledge Representation and Reasoning (KR'14)*, pages 52–61, 2014.

[12] S. Coste-Marquis, S. Konieczny, J.-G. Mailly, and P. Marquis. A translation-based approach for revision of argumentation frameworks. In *Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA'14)*, pages 77–85, 2014.

[13] S. Coste-Marquis, S. Konieczny, P. Marquis, and M.-A. Ouali. Selecting extensions in weighted argumentation frameworks. In *Proc. of the 4th International Conference on Computational Models of Argument, (COMMA'12)*, pages 342–349, 2012.

[14] S. Coste-Marquis, S. Konieczny, P. Marquis, and M.-A. Ouali. Weighted attacks in argumentation frameworks. In *Proc. of the 13th International Conference on the Principles of Knowledge Representation and Reasoning, (KR'12)*, pages 593–597, 2012.

[15] M. Cramer and M. Guillaume. Directionality of attacks in natural language argumentation. In *Proc. of the 4th Workshop on Bridging the Gap between Human and Automated Reasoning*, 2018. (in print).

[16] C. da Costa Pereira, A. Tettamanzi, and S. Villata. Changing one's mind: Erase or rewind? In *Proc. of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, pages 164–171, 2011.

[17] J. Delobelle, A. Haret, S. Konieczny, J. Mailly, J. Rossit, and S. Woltran. Merging of abstract argumentation frameworks. In *Proc. of the 15th International Conference the on Principles of Knowledge Representation and Reasoning, (KR'16)*, pages 33–42, 2016.

[18] J. Delobelle, S. Konieczny, and S. Vesic. On the aggregation of argumentation frameworks. In *Proc. of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 2911–2917, 2015.

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

[20] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2):457–486, 2011.

[21] P. E. Dunne, P. Marquis, and M. Wooldridge. Argument aggregation: Basic axioms and complexity results. In *Proc. of the 4th International Conference on Computational Models of Argument (COMMA'12)*, pages 129–140, 2012.

[22] D. M. Gabbay and O. Rodrigues. A numerical approach to the merging of argumentation networks. In *Proc. of the 13th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA'12)*, pages 195–212, 2012.

[23] S. Konieczny and R. Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

[24] G. H. Kramer. A dynamical model of political equilibrium. *Journal of Economic Theory*, 16(2):310–334, December 1977.

[25] D. G. Saari and V. R. Merlin. The copeland method 1; relationships and the dictionary. *Economic Theory*, 8, 1996.

[26] P. B. Simpson. On Defining Areas of Voter Choice: Professor Tullock on Stable Voting. *The Quarterly Journal of Economics*, 83(3):478–90, August 1969.

[27] F. A. Tohmé, G. A. Bodanza, and G. R. Simari. Aggregation of attack relations: A social-choice theoretical analysis of defeasibility criteria. In *Proc. of the 5th International Symposium on Foundations of Information and Knowledge Systems (FoIKS'08)*, pages 8–23, 2008.