

Merging of Abstract Argumentation Frameworks

Jérôme Delobelle

CRIL, CNRS
Université d’Artois, France
delobelle@cril.fr

Adrian Haret

Institute of Information Systems
TU Wien, Austria
haret@dbai.tuwien.ac.at

Sébastien Konieczny

CRIL, CNRS
Université d’Artois, France
konieczny@cril.fr

Jean-Guy Mailly

Institute of Information Systems
TU Wien, Austria
jmailly@dbai.tuwien.ac.at

Julien Rossit

LIPADE
Université Paris Descartes, France
julien.rossit@parisdescartes.fr

Stefan Woltran

Institute of Information Systems
TU Wien, Austria
woltran@dbai.tuwien.ac.at

Abstract

Formalizing dynamics of argumentation has received increasing attention over the last years. While AGM-like representation results for revision of argumentation frameworks (AFs) are now available, similar results for the problem of merging are still missing. In this paper, we close this gap and adapt model-based propositional belief merging to define extension-based merging operators for AFs. We state an axiomatic and a constructive characterization of merging operators through a family of rationality postulates and a representation theorem. Then we exhibit merging operators which satisfy the postulates. In contrast to the case of revision, we observe that obtaining a single framework as result of merging turns out to be a more subtle issue. Finally, we establish links between our new results and previous approaches to merging of AFs, which mainly relied on axioms from Social Choice Theory, but lacked AGM-like representation theorems.

Introduction

Over the last two decades argumentation has become a major research area in Knowledge Representation and Artificial Intelligence (Bench-Capon and Dunne 2007; Rahwan and Simari 2009). The work by Dung (1995) on abstract argumentation, in particular, is usually seen as a significant landmark in the consolidation of the field. Abstract argumentation is concerned with the evaluation of a set of arguments and their relations in order to extract subsets of the arguments—*extensions*—that can all be accepted together from some point of view. Dung’s argumentation frameworks (AFs), which are still the most widely used and investigated among the several argumentation formalisms, are directed graphs where nodes represent arguments and links correspond to one argument attacking another. The criteria or methods used to settle the acceptance of arguments are called “semantics” (see (Baroni, Caminada, and Giacomin 2011) for an overview).

Although substantial progress has been made on formalizing *dynamics* of argumentation frameworks (Baumann 2012; Bisquert et al. 2011; 2013; Boella, Kaci, and van der Torre 2009; Booth et al. 2013; Cayrol, Dupin de Saint

Cyr Bannay, and Lagasque-Schiex 2010; Doutre, Herzig, and Perrussel 2014; Kontarinis et al. 2013; Krümpelmann et al. 2012; Nouioua and Würbel 2014; Sakama 2014), only recently has this research direction matured towards representation theorems, a tool typically used in the AGM theory (Alchourrón, Gärdenfors, and Makinson 1985; Katsuno and Mendelzon 1991; Konieczny and Pérez 2002) of belief change. Here the goal is to formulate desired postulates for belief change that are then linked to all possible operators satisfying those postulates.

Such representation theorems for dynamics of argumentation have been first given in (Coste-Marquis et al. 2014) and later refined in (Diller et al. 2015), in the context of *revision* of argumentation frameworks. Both approaches work directly on the set of extensions of the involved AFs. To this end, some care is needed on the result of the revision, since it is not guaranteed that this result can be expressed via the extensions of a single AF. While Coste-Marquis et al. (2014) circumvent this problem by allowing multiple AFs as the result of the revision operator (interpreting the union of their extensions as the result), Diller et al. (2015) borrow concepts from Horn revision (Delgrande and Peppas 2015) and employ results on realizability of extension sets (Dunne et al. 2015) to tailor the operator such that the result is guaranteed to be represented by a single framework. A slightly different approach has been pursued in (Baumann and Brewka 2015), where a novel underlying monotonic logic was defined as basis for the revision task.

Thus, while AGM-style revision for argumentation frameworks is reasonably well understood, an application of the merging postulates due to Konieczny and Pérez (2002) is still open. In fact, previous work on aggregation of argumentation frameworks has focused on some kind of vote on the attack relations (Coste-Marquis et al. 2007; Tohmé, Bodanza, and Simari 2008; Delobelle, Konieczny, and Vesic 2015), with close links to properties from Social Choice Theory (Dunne, Marquis, and Wooldridge 2012). Properties proposed by Dunne et al. (2012) have been discussed by Delobelle et al. (2015). The conclusions are: first, few of them have to be taken as absolute requirements; second, there is room for two families of aggregation methods of argumentation frameworks: one focuses on the attack relations, the other focuses on the extensions. All previous approaches deal with the attack relation. In this work

we focus on the second family, i.e. we understand merging mainly through the changes it produces on the extensions of the AFs involved. Additionally, instead of focusing on properties coming from Social Choice Theory (and so dedicated to single sets of candidates), we rely on properties coming from propositional belief merging that seem more adequate when the set of "candidates" is structured. Indeed, a candidate here is not just any set of arguments, but has to satisfy the constraints given by an AF (just as a model has to satisfy the constraints given by the formulae). Finally, let us also mention (Caminada and Pigozzi 2011; Awad et al. 2014) which tackle a related but different problem. Here the question is to aggregate the different points of view (represented by the labellings) that some agents have about the same AF, while our approach allows each agent to have her own AF.

We believe that this form of *extension-based merging* is adequate for several application domains. Suppose, for instance, several expert groups debating a common topic. Each expert group delivers a set of arguments to be jointly accepted. In order to combine these proposals, the actual structure of the debate of each group is not central; rather, it is the sets of extensions we need to combine. Extension-based merging achieves precisely this. It is only in a second step that a possible structure of the debate is reconstructed, in terms of AFs that reflect the outcome of the merging.

In this paper we present the first representation theorem for extension-based argumentation merging. We first restate the merging postulates in the domain of abstract argumentation and derive a full representation theorem which is generic in the sense that it is not restricted to a particular argumentation semantics. We basically follow the multiple framework approach from (Coste-Marquis et al. 2014). However, we also discuss why working in the single framework approach from (Diller et al. 2015) is impossible for most semantics, in particular due to the fairness postulate. Finally, we establish links between our new results and previous approaches to merging of AFs.

Due to space constraints we omit proofs of our main results.

Background Notions

Abstract Argumentation

We start with a brief recapitulation of the necessary background on abstract argumentation due to Dung (1995).

Definition 1 (Dung 1995). *An abstract argumentation framework (AF) F is a directed graph $\langle A, R \rangle$ where A is a set of abstract entities called arguments and $R \subseteq A \times A$ is the attack relation between arguments. We use the notations $Arg(F) = A$ and $Att(F) = R$.*

The argument $a_i \in A$ attacks the argument $a_j \in A$ if $(a_i, a_j) \in R$. The set of arguments $S \subseteq A$ attacks a_j if $\exists a_i \in S$ such that a_i attacks a_j . a_i (respectively S) defends a_k against its attacker a_j if a_i (resp. S) attacks a_j .

We use \mathbf{AFs}_A to denote the set of all AFs built on the set of arguments A . Following Dung's approach, we define several types of *semantics*: these provide criteria for selecting

sets of arguments deemed to be acceptable with respect to a given AF.

Definition 2 (Dung 1995). *Given an AF $F = \langle A, R \rangle$, a set of arguments $S \subseteq A$ is conflict-free in F if $\forall a_i, a_j \in S, (a_i, a_j) \notin R$. A conflict-free set S is admissible in F if it defends all its arguments against each of their attackers. An admissible set S is:*

- a complete extension of F if each argument defended by S belongs to S ;
- a preferred extension of F if it is a \subseteq -maximal admissible set of F ;
- a stable extension of F if it attacks each argument in $A \setminus S$;
- the single grounded extension of F if it is the \subseteq -minimal complete extension of F .

The semantics are denoted respectively co, pr, st, gr . For each semantics σ , $Ext_\sigma(F)$ denotes the set of σ -extensions of F .

For an AF F with at least one extension, we say that an argument is sceptically accepted in F if it belongs to all of F 's extensions. An argument is credulously accepted in F if it belongs to at least one of F 's extensions. We denote by $sa_\sigma(F)$ (respectively $ca_\sigma(F)$) the set of sceptically (resp. credulously) accepted arguments in F .

AF Revision

This section recalls the AF revision setting proposed in (Coste-Marquis et al. 2014; Diller et al. 2015). Let us first introduce some terminology: we call a *candidate*¹ any set of arguments.

Definition 3 (Coste-Marquis et al. 2014). *Given $A = \{a_1, \dots, a_k\}$ a set of arguments, \mathbf{L}_A is the language generated by the following context-free grammar in BNF:*

$$\begin{aligned} arg & ::= a_1 | \dots | a_k \\ \Phi & ::= arg | \neg \Phi | (\Phi \wedge \Phi) | (\Phi \vee \Phi) \end{aligned}$$

For instance, $\varphi_1 = (a_1 \wedge a_2 \wedge a_3) \vee (a_1 \wedge \neg a_2 \wedge \neg a_3)$ expresses that in the revised epistemic state, a_1 must be accepted and a_2 and a_3 must be both accepted or both rejected. The epistemic status of such a formula φ from \mathbf{L}_A w.r.t. an argumentation framework $F \in \mathbf{AFs}_A$ for a semantics σ is given next.

Definition 4 (Coste-Marquis et al. 2014). *Let $c \subseteq A$ be a candidate and $\varphi \in \mathbf{L}_A$. The concept of satisfaction of φ by c , noted $c \models \varphi$, is defined inductively as follows:*

- If $\varphi = a \in A$, then $c \models \varphi$ if and only if $a \in c$,
- If $\varphi = (\varphi_1 \wedge \varphi_2)$, $c \models \varphi$ if and only if $c \models \varphi_1$ and $c \models \varphi_2$,
- If $\varphi = (\varphi_1 \vee \varphi_2)$, $c \models \varphi$ if and only if $c \models \varphi_1$ or $c \models \varphi_2$,
- If $\varphi = \neg \psi$, $c \models \varphi$ if and only if $c \not\models \psi$.

Then for any $F \in \mathbf{AFs}_A$, and any semantics σ , we say that:

¹Indeed, we expect these target candidates to become extensions at the end of the revision process, but a set of arguments cannot be called an extension if it is not associated with an AF (and a semantics); so before that they are *candidate to be an extension*.

- φ is (σ -)accepted with respect to F , noted $F \vdash_\sigma \varphi$, if $c \vdash \varphi$ for every $c \in \text{Ext}_\sigma(F)$,
- φ is (σ -)rejected with respect to F , noted $F \vdash_\sigma \neg\varphi$, if $c \vdash \varphi$ for no $c \in \text{Ext}_\sigma(F)$,
- φ is (σ -)undefined with respect to F in the remaining cases.

Inference \vdash_σ can be extended to the case of a set \mathcal{F} ; the question is to define what are the extensions of a set of AFs.

Definition 5. Given a semantics σ , we define the σ -extensions of a set of AFs \mathcal{F} as $\text{Ext}_\sigma(\mathcal{F}) = \bigcup_{F \in \mathcal{F}} \text{Ext}_\sigma(F)$. We can then extend \vdash_σ by saying that $\mathcal{F} \vdash_\sigma \varphi$ iff $c \vdash_\sigma \varphi$ for each $c \in \text{Ext}_\sigma(\mathcal{F})$.

Definition 6 (Coste-Marquis et al. 2014). Given a formula φ , \mathcal{A}_φ denotes the set of candidates satisfying φ . The formula φ is said to be consistent if and only if $\mathcal{A}_\varphi \neq \emptyset$.

Definition 7 (Coste-Marquis et al. 2014). A set \mathcal{C} of candidates is σ -representable if and only if there exists a set of AFs $\mathcal{S} \subseteq \text{AFs}_A$ such that $\mathcal{C} = \text{Ext}_\sigma(\mathcal{S})$.

Using σ -representability we define a notion of model which takes the semantics into account.

Definition 8 (Coste-Marquis et al. 2014). Given a formula $\varphi \in \mathbf{L}_A$ and a semantics σ , the set of models of φ is defined as follows:

$$\mathcal{A}_\varphi^\sigma = \{c \in \mathcal{A}_\varphi \mid \{c\} \text{ is } \sigma\text{-representable}\}.$$

A formula $\varphi \in \mathbf{L}_A$ is σ -representable if and only if $\mathcal{A}_\varphi^\sigma$ is σ -representable. We use $\varphi \equiv_\sigma \psi$ to denote $\mathcal{A}_\varphi^\sigma = \mathcal{A}_\psi^\sigma$.

It is now possible to define revision operators for AFs. In (Coste-Marquis et al. 2014), AF revision operators are defined such that their result is a set of AFs whose extensions satisfy rationality postulates in the style of (Katsuno and Mendelzon 1991). AF revision operators are constructed by a two step-process: the first step is to select the set of revised candidates; the second step is to generate the AFs which represent these candidates.

Definition 9 (Coste-Marquis et al. 2014). For a given semantics σ , an AF revision operator is a mapping from an AF $F \in \text{AFs}_A$ and a formula $\varphi \in \mathbf{L}_A$ to a set of AFs denoted $F \star \varphi$ such that $\text{Ext}_\sigma(F \star \varphi) = \min(\mathcal{A}_\varphi^\sigma, \leq_F)$, with \leq_F a pre-order between candidates.

It has been proven that if \leq_F satisfies the properties of a faithful assignment, then \star satisfies the rationality postulates adapted from (Katsuno and Mendelzon 1991).

Definition 10 (Coste-Marquis et al. 2014). For a given semantics σ , a faithful assignment is a mapping from an AF F to a total pre-order between candidates \leq_F such that:

- if $c_1 \in \text{Ext}_\sigma(F)$ and $c_2 \in \text{Ext}_\sigma(F)$, then $c_1 \simeq_F c_2$;
- if $c_1 \in \text{Ext}_\sigma(F)$ and $c_2 \notin \text{Ext}_\sigma(F)$, then $c_1 <_F c_2$.

The second step of the process is the generation of the revised AFs.

Definition 11 (Coste-Marquis et al. 2014). Given a semantics σ , a generation operator \mathcal{AF}_σ is a mapping from a set of candidates \mathcal{C} to a set of AFs such that $\text{Ext}_\sigma(\mathcal{AF}_\sigma(\mathcal{C})) = \mathcal{C}$. Given a faithful assignment which maps every AF F to a

total pre-order \leq_F , a semantics σ and a generation operator \mathcal{AF}_σ , we define the corresponding revision operator as follows. For each AF F and each formula φ ,

$$F \star \varphi = \mathcal{AF}_\sigma(\min(\mathcal{A}_\varphi^\sigma, \leq_F))$$

Several generation operators are proposed to ensure the result satisfies additional desirable properties, such as minimal change of the attack relation, or minimal cardinality of the result.

Throughout this process, the working assumption is that the result of AF revision is a set of AFs. Having a set of AFs as output reflects the uncertainty of the result. But for certain applications it can be required that the agent obtains a single AF from the revision process. This cannot be guaranteed with our working definition of AF revision, since the set of candidates which are selected at the first step of the revision is in general not *realizable* (Dunne et al. 2014), i.e. it is not possible to find a single AF such that its σ -extensions are the expected candidates. To overcome this issue, (Diller et al. 2015) propose a new approach which ensures that the selected candidates are realizable with respect to the chosen semantics.

Definition 12 (Diller et al. 2015). Given a semantics σ and a pre-order \leq , \leq is σ -compliant if for every consistent formula $\varphi \in \mathbf{L}_A$ it holds that $\min(\mathcal{A}_\varphi^\sigma, \leq)$ is σ -realizable.

Using a faithful assignment where all the pre-orders \leq_F are σ -compliant, it is now possible to define a revision operator its output is a single AF.

Merging Operators for AFs

In this section, we propose an adaptation of the rationality postulates for logical belief merging in the setting of abstract argumentation, and we give a constructive characterization of the operators which satisfy these postulates, before presenting a particular family of operators satisfying them.

In the following, $\mathcal{F} = (F_1, \dots, F_k)$ denotes the profile² of AFs representing the group of agents, with each AF built on the set of arguments A , and so each formula used in the following postulates belongs to the language \mathbf{L}_A associated with A .

The general definition of a merging operator is as follows.

Definition 13. An AF merging operator is a mapping from a profile \mathcal{F} and a formula $\mu \in \mathbf{L}_A$ to a set of AFs $\Delta_\mu(\mathcal{F})$.

The formula μ is called the integrity constraint, and it represents some information that must hold in the result of the merging. It can be some physical law, some legal law, etc.

We are interested in a more specific class of merging operators. Let us first introduce some concepts.

We say that two profiles \mathcal{F}_1 and \mathcal{F}_2 are equivalent, denoted $\mathcal{F}_1 \equiv \mathcal{F}_2$, if and only if there exists a bijection ρ from \mathcal{F}_1 to \mathcal{F}_2 such that $\forall F \in \mathcal{F}_1, \rho(F) = F$. We define the conjunction of a profile $\mathcal{F} = (F_1, \dots, F_k)$ as the set of extensions that are in all F_i : $\text{Ext}_\sigma(\bigwedge \mathcal{F}) = \bigcap_{F \in \mathcal{F}} \text{Ext}_\sigma(F)$.

²A profile is a tuple. The union of two profiles is just the concatenation of the two tuples.

Definition 14. For a given semantics σ and an AF merging operator Δ , we say that Δ is an AF IC-merging operator if, for any profile \mathcal{F} and formula $\mu \in \mathbf{L}_A$, Δ satisfies the following postulates:

- (M0) $Ext_\sigma(\Delta_\mu(\mathcal{F})) \subseteq \mathcal{A}_\mu^\sigma$
- (M1) If $\mathcal{A}_\mu^\sigma \neq \emptyset$, then $Ext_\sigma(\Delta_\mu(\mathcal{F})) \neq \emptyset$
- (M2) If $Ext_\sigma(\bigwedge \mathcal{F}) \cap \mathcal{A}_\mu^\sigma \neq \emptyset$, then $Ext_\sigma(\Delta_\mu(\mathcal{F})) = Ext_\sigma(\bigwedge \mathcal{F}) \cap \mathcal{A}_\mu^\sigma$
- (M3) If $\mathcal{F}_1 \equiv \mathcal{F}_2$ and $\mu_1 \equiv_\sigma \mu_2$, then $Ext_\sigma(\Delta_{\mu_1}(\mathcal{F}_1)) = Ext_\sigma(\Delta_{\mu_2}(\mathcal{F}_2))$
- (M4) If $Ext_\sigma(F_1) \subseteq \mathcal{A}_\mu^\sigma$ and $Ext_\sigma(F_2) \subseteq \mathcal{A}_\mu^\sigma$, then $Ext_\sigma(\Delta_\mu((F_1, F_2))) \cap Ext_\sigma(F_1) \neq \emptyset$ implies $Ext_\sigma(\Delta_\mu((F_1, F_2))) \cap Ext_\sigma(F_2) \neq \emptyset$
- (M5) $Ext_\sigma(\Delta_\mu(\mathcal{F}_1)) \cap Ext_\sigma(\Delta_\mu(\mathcal{F}_2)) \subseteq Ext_\sigma(\Delta_\mu(\mathcal{F}_1 \cup \mathcal{F}_2))$
- (M6) If $Ext_\sigma(\Delta_\mu(\mathcal{F}_1)) \cap Ext_\sigma(\Delta_\mu(\mathcal{F}_2)) \neq \emptyset$, then $Ext_\sigma(\Delta_\mu(\mathcal{F}_1 \cup \mathcal{F}_2)) \subseteq Ext_\sigma(\Delta_\mu(\mathcal{F}_1)) \cap Ext_\sigma(\Delta_\mu(\mathcal{F}_2))$
- (M7) $Ext_\sigma(\Delta_{\mu_1}(\mathcal{F})) \cap \mathcal{A}_{\mu_2}^\sigma \subseteq Ext_\sigma(\Delta_{\mu_1 \wedge \mu_2}(\mathcal{F}))$
- (M8) If $Ext_\sigma(\Delta_{\mu_1}(\mathcal{F})) \cap \mathcal{A}_{\mu_2}^\sigma \neq \emptyset$, then $Ext_\sigma(\Delta_{\mu_1 \wedge \mu_2}(\mathcal{F})) \subseteq Ext_\sigma(\Delta_{\mu_1}(\mathcal{F})) \cap \mathcal{A}_{\mu_2}^\sigma$

These postulates are direct translations of the ones from (Konieczny and Pérez 2002) for the expected extensions of the merging of a profile of argumentation frameworks. Here is their interpretation in this setting.

(M0) ensures that the selected extensions satisfy the constraints. (M1) says that as soon as the constraints are σ -consistent (i.e. at least one set of extensions can satisfy them), then the result of the merging will be non trivial (i.e. there will be at least one selected extension). (M2) ensures that if there is a consensus among the AFs of the profile on some extensions, then the result of the merging will be exactly these extensions. (M3) says that the result of the merging does not depend on the order in which AFs are arranged in the profile and on the syntax of the constraints. (M4) is a fairness postulate, saying that all AFs have the same impact. In particular, if we merge only two AFs and if the result is compatible with the extensions of one of the AFs, then it also has to be compatible with the other one. This ensures that the result can not be biased towards one of the two AFs. (M5) and (M6) deal with what happens when we do the union of two profiles. (M5) says that if there are extensions that are independently chosen as result of the merging for two profiles \mathcal{F}_1 and \mathcal{F}_2 , then they also have to be result of the merging if we join the two profiles. And (M6) adds that in this case they are the only extensions in the merging the joined profile. (M7) and (M8) give the expected behaviour of the operators when a profile is merged with a conjunction as constraint. This expresses the fact that good extensions with a given constraint μ_1 are still good extensions if we just (consistently) strengthen the constraint by μ_2 .

Similar to propositional merging operators (Konieczny and Pérez 2002), AF merging operators can be represented by plausibility pre-orders. In our case these pre-orders are used to represent the respective plausibility of candidates given a profile:

Definition 15. For a given semantics σ , a syncretic assignment maps any profile \mathcal{F} to a total pre-order on candidates $\leq_{\mathcal{F}}$ such that, for all $\mathcal{F}, \mathcal{F}_1, \mathcal{F}_2$ and for all F_1, F_2 , the following properties are satisfied:

1. If $c_1 \in Ext_\sigma(\bigwedge \mathcal{F})$, $c_2 \in Ext_\sigma(\bigwedge \mathcal{F})$, then $c_1 \simeq_{\mathcal{F}} c_2$
2. If $c_1 \in Ext_\sigma(\bigwedge \mathcal{F})$, $c_2 \notin Ext_\sigma(\bigwedge \mathcal{F})$, then $c_1 <_{\mathcal{F}} c_2$
3. $\forall c_1 \in Ext_\sigma(F_1), \exists c_2 \in Ext_\sigma(F_2)$ s.t. $c_2 \leq_{(F_1, F_2)} c_1$
4. If $c_1 \leq_{\mathcal{F}_1} c_2$ and $c_1 \leq_{\mathcal{F}_2} c_2$, then $c_1 \leq_{\mathcal{F}_1 \cup \mathcal{F}_2} c_2$
5. If $c_1 <_{\mathcal{F}_1} c_2$ and $c_1 \leq_{\mathcal{F}_2} c_2$, then $c_1 <_{\mathcal{F}_1 \cup \mathcal{F}_2} c_2$

Syncretic assignments give us a way to characterize AF IC-merging operators.

Proposition 1. Given a semantics σ , an operator Δ satisfies (M0)-(M8) iff there exists a syncretic assignment that maps each profile \mathcal{F} to a total pre-order $\leq_{\mathcal{F}}$ between candidates such that $Ext_\sigma(\Delta_\mu(\mathcal{F})) = \min(\mathcal{A}_\mu^\sigma, \leq_{\mathcal{F}})$.

Before moving to the second step of the process, let us notice that (as in the propositional case) merging operators are a generalization of revision operators.

Proposition 2. Given a merging operator Δ , we define \star as follows: $F \star \mu = \Delta_\mu((F))$. If Δ satisfies (M0)-(M8), then \star satisfies the revision postulates (AE1)-(AE6) proposed in (Coste-Marquis et al. 2014).

Similarly to the revision approach defined in (Coste-Marquis et al. 2014), our approach for merging AFs is a two-step process: first the computation of the expected extensions, then the generation of the corresponding AFs. In the next section, we define generation operators which are a generalization of those proposed in (Coste-Marquis et al. 2014). This implies that the revision operators defined in Proposition 2 correspond to the existing ones, provided that the generation step is performed as described below.

Generation of the Merged AFs

Here we study the generation of the AFs which result from the merging process. This is a generalisation of the generation operators defined in (Coste-Marquis et al. 2014).

Definition 16. Given a semantics σ , a generation operator \mathcal{AF}_σ is a mapping from a set of candidates \mathcal{C} and a profile \mathcal{F} to a set of AFs such that $Ext_\sigma(\mathcal{AF}_\sigma(\mathcal{C}, \mathcal{F})) = \mathcal{C}$.

We notice that the profile is not used in the general definition of a generation operator, but it will be useful to define some specific ones. Now we give the full definition of our merging operators, based on the selection of desired set of extensions, and on the generation of the corresponding AFs:

Definition 17. Given a semantics σ , a syncretic assignment which maps every profile \mathcal{F} to a total pre-order $\leq_{\mathcal{F}}$, and a generation operator \mathcal{AF}_σ , we define a merging operator Δ as follows. For each formula μ we say that $\Delta_\mu(\mathcal{F}) = \mathcal{AF}_\sigma(\min(\mathcal{A}_\mu^\sigma, \leq_{\mathcal{F}}), \mathcal{F})$.

Proposition 3. Every merging operator Δ defined following Definition 17 satisfies the postulates (M0)-(M8).

Let us now define specific generation operators for the merging process. Indeed, the existing generation operator

proposed for revision combine minimal change on the attack relation and minimal cardinality of the result (Coste-Marquis et al. 2014). But the notion of minimal change on the attack relation in the revision case only considers a single input AF. Here we have as input a full profile of AFs, so we can also take this information into account for defining better generation operators. Hence we will need two aggregation functions, the first one to compute the distance of each generated AF with respect to the initial AFs of the input, and the second one to compute the distance of a (whole) set of generated AFs.

Definition 18. An aggregation function is a function \otimes associating a non-negative real to every finite tuple of non-negative real and satisfying the following properties:

- If $x \leq y$, then $\otimes(x_1, \dots, x, \dots, x_n) \leq \otimes(x_1, \dots, y, \dots, x_n)$. (non-decreasingness)
- $\otimes(x_1, \dots, x_n) = 0$ iff $x_1 = \dots = x_n = 0$. (minimality)
- For every non-negative real x , $\otimes(x) = x$. (identity)

The first approach gives the priority to minimal distance, and then to the output set cardinality:

Definition 19. Given a semantics σ . Let \mathcal{C} be a set of candidates, dg be a (pseudo-)distance between graphs, \mathcal{F} be a profile of AFs and \otimes, \odot be two aggregation functions, $\mathcal{AF}_\sigma^{dg \otimes, \odot}$ is defined as:

$$\mathcal{AF}_\sigma^{dg \otimes, \odot}(\mathcal{C}, \mathcal{F}) = \gamma\{\mathcal{F}' \in \text{sets} \mid |\mathcal{F}'| \text{ is minimal}\}$$

with

$$\text{sets} = \{\mathcal{F}' \mid \text{Ext}_\sigma(\mathcal{F}') = \mathcal{C} \text{ and } \otimes_{F \in \mathcal{F}} \odot_{F' \in \mathcal{F}'} dg(F, F') \text{ is minimal}\}$$

and γ a selection function which maps a set of sets of AFs $\mathcal{F}s = \{\mathcal{F}_1, \dots, \mathcal{F}_k\}$ to $\mathcal{F}_i \in \mathcal{F}s$.

γ is a tie-break function, used to select only one of the possible minimal sets if there are several ones.

A second approach consists in giving priority to the minimality of the output cardinality. It builds first sets of systems that cover the set of candidates with a minimal number of systems, and then chooses the sets which minimize the change on the attack relation.

Definition 20. Given a semantics σ . Let \mathcal{C} be set of candidates, dg be a (pseudo-)distance between graphs, \mathcal{F} be a profile of AFs and \otimes, \odot be two aggregation functions, $\mathcal{AF}_\sigma^{card \otimes, \odot}$ is defined by:

$$\mathcal{AF}_\sigma^{card \otimes, \odot}(\mathcal{C}, \mathcal{F}) = \gamma\{\mathcal{F}' \in \text{sets} \mid \otimes_{F \in \mathcal{F}} \odot_{F' \in \mathcal{F}'} dg(F, F') \text{ is minimal}\}$$

with $\text{sets} = \{\mathcal{F}' \mid \text{Ext}_\sigma(\mathcal{F}') = \mathcal{C} \text{ and } |\mathcal{F}'| \text{ is minimal}\}$.

Let us illustrate these two approaches for generation functions on an example.

Example 1. We illustrate the behaviour of both $\mathcal{AF}_\sigma^{dg \otimes, \odot}$ and $\mathcal{AF}_\sigma^{card \otimes, \odot}$ with the aggregation functions $\otimes = \text{avg}$ (the standard arithmetic average), and $\odot = \Sigma$. We consider the profile $\mathcal{F} = (F_1, F_2, F_3)$ given at Figure 1. We suppose that the first step has conducted to the set of merged candidates $\{\{a_1, a_2, a_4\}, \{a_2, a_3, a_4\}\}$. With the generation operator $\mathcal{AF}_\sigma^{dg_{\text{avg}}, \Sigma}$, we obtain the set of AFs $\mathcal{F}'_1 = \{F'_1, F'_2\}$




Figure 1: Three AFs to be Merged




Figure 2: Result of Generation with $\mathcal{AF}_\sigma^{dg_{\text{avg}}, \Sigma}$

described at Figure 2. We remark that each of them corresponds to one of the candidates $\{\{a_1, a_2, a_4\}, \{a_2, a_3, a_4\}\}$. The distances between F'_1 and the AFs F_1, F_2, F_3 are respectively 4, 3, 0 so the distance between F'_1 and \mathcal{F} is 7. Similarly we obtain the distance between F'_2 and \mathcal{F} which is 5. So the average distance between \mathcal{F}'_1 and \mathcal{F} is 6. This result is minimal with respect to the distance between AFs, but not with respect to the cardinality.

On the opposite, we obtain a singleton $\mathcal{F}'_2 = \{F'_3\}$ as described at Figure 3 when the generation is performed by $\mathcal{AF}_\sigma^{card_{\text{avg}}, \Sigma}$. Here the result is obviously minimal with re-




Figure 3: Result of Generation with $\mathcal{AF}_\sigma^{card_{\text{avg}}, \Sigma}$

spect to the cardinality. However, compared to \mathcal{F}'_1 it is not minimal with respect to the distance between AFs: the distances between F'_3 and the AFs F_1, F_2, F_3 are respectively 4, 5, 2, so the distance between \mathcal{F}'_2 and \mathcal{F} is 11.

Distance-based merging operators

We know that merging operators can be represented by pre-orders between candidates. One possible way to build these pre-orders is to use distance-based operators. So let us now define distance-based merging operators. This definition requires the introduction of the concept of (pseudo-)distance.

Definition 21. A (pseudo-)distance d between sets of arguments is a mapping from two sets of arguments to a non-negative real number such that for each c_1, c_2 two sets of arguments,

- $d(c_1, c_2) = 0$ iff $c_1 = c_2$;
- $d(c_1, c_2) = d(c_2, c_1)$.

The distance between a candidate c_1 and a set of candidates C is defined by $d(c_1, C) = \min_{c_2 \in C} (d(c_1, c_2))$.

Definition 22 (Distance-based merging operator). *Given a semantics σ . Let d be a distance between candidates, \otimes be an aggregation functions, and \mathcal{AF}_σ be a generation operator. We define the pre-order $\leq_{\mathcal{F}}^{\otimes, d}$ between candidates as follows:*

$$c_1 \leq_{\mathcal{F}}^{\otimes, d} c_2 \text{ iff } \bigotimes_{F \in \mathcal{F}} d(c_1, \text{Ext}_\sigma(F)) \leq \bigotimes_{F \in \mathcal{F}} d(c_2, \text{Ext}_\sigma(F))$$

Now, the merging operator $\Delta^{\otimes, d, \mathcal{AF}_\sigma}$ is defined by

$$\Delta_\mu^{\otimes, d, \mathcal{AF}_\sigma}(\mathcal{F}) = \mathcal{AF}_\sigma(\min(\mathcal{A}_\mu^\sigma, \leq_{\mathcal{F}}^{\otimes, d}), \mathcal{F})$$

We need some additional properties to characterize the aggregation functions which lead to distance-based operators which satisfy the postulates:

- For any permutation π , $\otimes(x_1, \dots, x_n) = \otimes(\pi(x_1, \dots, x_n))$. (symmetry)
- If $\otimes(x_1, \dots, x_n) \leq \otimes(y_1, \dots, y_n)$, then $\otimes(x_1, \dots, x_n, z) \leq \otimes(y_1, \dots, y_n, z)$. (composition)
- If $\otimes(x_1, \dots, x_n, z) \leq \otimes(y_1, \dots, y_n, z)$, then $\otimes(x_1, \dots, x_n) \leq \otimes(y_1, \dots, y_n)$. (decomposition)

Proposition 4. *For any semantics σ , and any pseudo-distance between candidates d , any generation operator \mathcal{AF}_σ and any aggregation function \otimes which satisfies symmetry, composition and decomposition, the merging operator $\Delta^{\otimes, d, \mathcal{AF}_\sigma}$ satisfies the postulates (M0)–(M8).*

Now, let us illustrate our approach with a concrete operator from this family.

Example 2. *We consider the Hamming distance between candidates: $d_H(c_1, c_2) = |(c_1 \setminus c_2) \cup (c_2 \setminus c_1)|$. We want to merge the profile $\mathcal{F} = (F_1, F_2, F_3)$ described at Figure 1, under the stable semantics, with the generation operator $\mathcal{AF}_{\sigma, \mathcal{F}}^{\text{card}_{\text{avg}, \Sigma}}$, and with the integrity constraint $\mu = a_2 \wedge a_4 \wedge (a_1 \vee a_3)$, which expresses that for each extension, the arguments a_2 and a_4 must belong to it, and at least one of the arguments a_1 and a_3 . So we want to compute $\Delta_\mu^{\Sigma, d_H, \mathcal{AF}_{\sigma, \mathcal{F}}^{\text{card}_{\text{avg}, \Sigma}}}(\mathcal{F})$. We also compute the result of the merging when the generation operator is $\mathcal{AF}_{\sigma, \mathcal{F}}^{\text{dg}_{\text{avg}, \Sigma}}$, which leads to $\Delta_\mu^{\Sigma, d_H, \mathcal{AF}_{\sigma, \mathcal{F}}^{\text{dg}_{\text{avg}, \Sigma}}}(\mathcal{F})$.*

The stable extensions of the AFs are $\text{Ext}_{st}(F_1) = \{\{a_1, a_3, a_4\}, \{a_2, a_3, a_4\}\}$, $\text{Ext}_{st}(F_2) = \{\{a_2, a_4\}\}$ and $\text{Ext}_{st}(F_3) = \{\{a_1, a_2, a_4\}\}$. The following table exhibits the distance between each model of μ and the extensions of F_1 , F_2 and F_3 .

μ	F_1 $\{a_1, a_3, a_4\}$ $\{a_2, a_3, a_4\}$	F_2 $\{a_2, a_4\}$	F_3 $\{a_1, a_2, a_4\}$	Σ
$\{a_1, a_2, a_4\}$	2	1	0	3
$\{a_2, a_3, a_4\}$	0	1	2	3
$\{a_1, a_2, a_3, a_4\}$	1	2	1	4

Now the distance between each model c of μ and the profile \mathcal{F} is the sum of the distances between c and each of the AFs in \mathcal{F} . We obtain the following distances: $d_{\Sigma, d_H}(\{a_1, a_2, a_4\}, \mathcal{F}) = 2 + 1 + 0 = 3$, $d_{\Sigma, d_H}(\{a_2, a_3, a_4\}, \mathcal{F}) = 0 + 1 + 2 = 3$ and $d_{\Sigma, d_H}(\{a_2, a_3, a_3, a_4\}, \mathcal{F}) = 1 + 2 + 1 = 4$.

So the candidates which are selected at the merging step are $\{\{a_1, a_2, a_4\}, \{a_2, a_3, a_4\}\}$. We observe that this result conducts to the generation step which is described at Example 1. So $\Delta_\mu^{\Sigma, d_H, \mathcal{AF}_{\sigma, \mathcal{F}}^{\text{dg}_{\text{avg}, \Sigma}}}(\mathcal{F}) = \mathcal{F}'_1$ described at Figure 2, and $\Delta_\mu^{\Sigma, d_H, \mathcal{AF}_{\sigma, \mathcal{F}}^{\text{card}_{\text{avg}, \Sigma}}}(\mathcal{F}) = \mathcal{F}'_2$ given at Figure 3.

We have shown that it is possible to satisfy all the postulates together. But let us notice now that it is not that an easy task, and let us illustrate this by checking what are the postulates satisfied by the aggregation operators introduced in (Delobelle, Konieczny, and Vesic 2015).

Roughly speaking, these operators compute the result of merging within two steps: they first aggregate a set of AFs into a WAF (Weighted Aggregation Framework (Dunne et al. 2011; Coste-Marquis et al. 2012a; 2012b)) by counting how many AFs agree with each attack; then best ordering rules and relaxing extensions techniques are used to select extensions produced by this WAF to provide the result of merging. We do not have space here to recall the full definitions of these operators, so please look at (Delobelle, Konieczny, and Vesic 2015).

One can notice that these WAF based approaches do not allow to express any integrity constraints μ . Thus, in what follows, we focus on the cases where μ is formula representing a tautology, i.e. the whole set of arguments expressed among each AFs to merge. Note that we drop the consideration of (M0), (M7) and (M8) which are trivially satisfied when $\mu \equiv \top$ (resp. $\mu_1 \equiv \mu_2 \equiv \top$), namely when μ (resp. μ_1 and μ_2) do not constrain the result of merging.

Proposition 5. *Let μ be a formula such that $\mu \equiv \top$. Given a semantics $\sigma \in \{co, pr, st, gr\}$:*

- $FUS_{All}^{\sigma, \text{best}_i^\oplus}$ satisfies (M3) but violates (M1), (M2), and (M4)–(M6);
- $FUS_{AllNT}^{\sigma, \text{best}_i^{\oplus 1, \oplus 2}}$ and $FUS_{MajNT}^{\sigma, \text{best}_i^{\oplus 1, \oplus 2}}$ satisfy (M1) and (M3), but violate (M2), and (M4)–(M6).

So one can remark that these operators satisfy almost none of the IC merging postulates, in particular (M4)–(M6), that are clearly aggregation properties, are not satisfied.

Instantiating the Result as a Single AF

Our working definition of merging operators assumes that the output of merging is a set of AFs. As for revision (Coste-Marquis et al. 2014), this is prompted by the fact that (unlike for logical approaches (Alchourrón, Gärdenfors, and Makinson 1985; Katsuno and Mendelzon 1991; Konieczny and Pérez 2002; Delgrande and Peppas 2015)) there is no disjunction in the language allowing us to combine the possible results³. In this section, however, we consider the possibility of instantiating the output of merging as a *single* AF.

Concretely, we want to see whether there exist merging operators Δ which (i) satisfy postulates (M0) – (M8), (ii)

³Note that for partial meet contraction and revision functions the result is also a set (of maximal consistent subsets). But in this setting they can be combined using intersection.

are constructed using distance-based methods (see Definition 22), and (iii) have a single AF as output. The formal definitions are given below.

Definition 23. For a given semantics σ , a resolute AF merging operator is a mapping from a set of AFs \mathcal{F} and a formula $\mu \in \mathbf{L}_A$ to an AF $\Delta_\mu(\mathcal{F})$. A resolute AF merging operator is a resolute IC-merging operator if it satisfies **(M0)** – **(M8)**.

Note that when dealing with resolute merging operators, $Ext_\sigma(\Delta_\mu(\mathcal{F}))$ refers to the set of extensions of a single AF rather than the union of the sets of extensions of several AFs. The notion of a syncretic assignment (Definition 15) stays unchanged, as it does not depend on how the output of a merging operator is defined. Hence, given a set \mathcal{F} of AFs, we can try to apply familiar distance-based notions of Definition 22 to generate a set of extensions as our merging output.

However, our new setting comes with an extra layer of difficulty: if we want to use distance-based methods for defining resolute merging operators, then we need the resulting set of candidates to be the σ -extensions of some *single* AF. In other words, we require that for any set of AFs \mathcal{F} , and for any formula $\mu \in \mathbf{L}_A$, $\min(\mathcal{A}_\mu^\sigma, \leq_{\mathcal{F}}^{\otimes, d})$ is σ -realizable, or that $\leq_{\mathcal{F}}^{\otimes, d}$ is σ -compliant (see Definition 12). But as Example 3 below shows, familiar distance-based methods have a problem in this respect.⁴

Example 3. We want to merge F_1 and F_2 from Example 2 and constraint $\mu = (a_1 \oplus a_4) \wedge (a_2 \oplus a_4)$, using Hamming distance and under any semantics $\sigma \in \{pr, st\}$. We have $\mathcal{F}' = (F_1, F_2)$ and $Ext_\sigma(F_1) = \{\{a_1, a_3, a_4\}, \{a_2, a_3, a_4\}\}$, $Ext_\sigma(F_2) = \{\{a_2, a_4\}\}$, $\mathcal{A}_\mu^\sigma = \{\{a_1, a_2\}, \{a_1, a_2, a_3\}, \{a_4\}, \{a_3, a_4\}\}$. We obtain the following distances: $d_\Sigma(\{a_1, a_2\}, \mathcal{F}') = 3 + 2 = 5$, $d_\Sigma(\{a_1, a_2, a_3\}, \mathcal{F}') = 2 + 3 = 5$, $d_\Sigma(\{a_4\}, \mathcal{F}') = 2 + 1 = 3$, $d_\Sigma(\{a_3, a_4\}, \mathcal{F}') = 1 + 2 = 3$. Thus at the merging step the selected candidates are $\{a_4\}$ and $\{a_3, a_4\}$. But we know that for any of the semantics considered, the set of extensions of any AF must be incomparable. So there is no F such that $Ext_\sigma(F) = \{\{a_4\}, \{a_3, a_4\}\}$ (see (Dunne et al. 2015)).

Thus, pre-orders constructed with Hamming distance are not necessarily *st*- or *pr*-compliant. The σ -compliance condition is both a blessing and a curse. It is a curse because it renders familiar distance-based methods ineffectual. It is a blessing because with it in place we have the following representation result.

Proposition 6. A resolute merging operator Δ satisfies postulates **(M0)**–**(M8)** if and only if there exists a syncretic assignment that maps each profile \mathcal{F} to a σ -compliant total pre-order $\leq_{\mathcal{F}}$ such that $Ext_\sigma(\Delta_\mu(\mathcal{F})) = \min(\mathcal{A}_\mu, \leq_{\mathcal{F}})$.

With Proposition 6 we know that resolute IC-merging operators are characterized by syncretic assignments where the pre-orders are σ -compliant, and hence nurture the hope that we can perhaps find some other distance that delivers σ -compliant pre-orders. Unfortunately, this turns out not to be the case, as Proposition 7 below shows.

⁴This situation is reminiscent of the problem faced when doing propositional belief change in certain fragments, such as Horn (Delgrande and Peppas 2015; Haret, Rümmele, and Woltran 2015).

Proposition 7. There are no resolute IC-merging operators for stable (resp. preferred, grounded, complete) semantics.

Proof. Assume there exists a resolute IC-merging operator Δ . We first look at semantics $\sigma \in \{st, pr, gr\}$. Take F_7 and F_8 , two AFs as in Figure 4 below, and the constraint $\mu = a$, with $A_\mu = \{\{a\}, \{a, b\}\}$.




Figure 4

The σ -extensions of F_7 and F_8 are $Ext_\sigma(F_7) = \{\{a\}\}$, $Ext_\sigma(F_8) = \{\{a, b\}\}$. Since Δ is a resolute IC-merging operator it satisfies postulates **(M0)** – **(M8)**. **(M0)** and **(M1)** require that $Ext_\sigma(\Delta_\mu((F_7, F_8)))$ is a non-empty subset of \mathcal{A}_μ^σ , while **(M4)** requires that both $\{a\}$ and $\{a, b\}$ are elements of $Ext_\sigma(\Delta_\mu((F_7, F_8)))$. Thus the only possibility is that $Ext_\sigma(\Delta_\mu((F_7, F_8))) = \{\{a\}, \{a, b\}\}$. But we know that the set of grounded extensions of any AF must contain exactly one element, and that when σ is preferred or stable semantics, the set of σ -extensions of any AF must be incomparable. The set $\{\{a\}, \{a, b\}\}$ satisfies neither of these properties, hence there is no single AF F that represents $\{\{a\}, \{a, b\}\}$ under semantics σ . This contradicts our assumption that Δ is a resolute IC-merging operator.

For complete semantics, consider the profile (F_7, F_9) , where F_7 and F_9 are as in Figure 4, and $\mu = a \oplus b$. Then $Ext_{co}(F_7) = \{\{a\}\}$ and $Ext_{co}(F_9) = \{\{b\}\}$, and applying **(M0)**, **(M1)** and **(M4)** as before, we get that $Ext_\sigma(\Delta_\mu((F_7, F_9))) = \{\{a\}, \{b\}\}$. However, the set $\{\{a\}, \{b\}\}$ is not *co*-realizable, as it does not contain the intersection of all its elements (see (Dunne et al. 2015)). \square

Proposition 7 shows that there are no resolute merging operators for many of the standard semantics used in argumentation. One way to circumvent this situation is to restrict μ in a suitable way, though an easy solution is not obvious and we reserve this line of research for future work. The immediate import of our result, however, is that simply restricting the output of AF merging operators to a single AF does not work. This lends support to our original option of representing the merging output with a set of AFs.

Merging Operators and Aggregation Axioms

In this section, we investigate how our merging operators relate to the properties proposed by Dunne et al. (2012) for the aggregation of profiles of argumentation frameworks.

We recall that \mathbf{AFs}_A denotes the set of all argumentation frameworks defined from a (finite) set of arguments A known by all the agents, denoted by \mathcal{N} . An aggregation function γ is defined by $\gamma : \mathbf{AFs}_A^n \rightarrow \mathbf{AFs}_A$. A profile $\mathcal{F} = (F_1, \dots, F_n)$ is a tuple in \mathbf{AFs}_A^n . Unless stated explicitly all the properties are defined $\forall \mathcal{F} \in \mathbf{AFs}_A^n$.

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

(ANON) $\forall \mathcal{F}' \in \Pi(\mathcal{F}) : \gamma(\mathcal{F}) = \gamma(\mathcal{F}')$

Non-Triviality. An argumentation framework is non-trivial, for a semantics σ , if it has at least one non-empty extension: $|Ext_\sigma(F)| \geq 1$ and $Ext_\sigma(F) \neq \{\emptyset\}$. Let us note $\mathbf{AFs}_A^{NT\sigma}$ the set of non-trivial (for the semantics σ) argumentation frameworks. The aggregation function γ is σ -**strongly non-trivial** if the output is always non-trivial :

(σ -SNT) $\gamma(\mathcal{F}) \in \mathbf{AFs}_A^{NT\sigma}$

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

(σ -WNT) $\forall \mathcal{F} \in (\mathbf{AFs}_A^{NT\sigma})^n : \gamma(\mathcal{F}) \in \mathbf{AFs}_A^{NT\sigma}$

Decisiveness. An argumentation framework is decisive, for a semantics σ , if it has exactly one non-empty extension: $|Ext_\sigma(F)| = 1$ and $Ext_\sigma(F) \neq \{\emptyset\}$. Let us note $\mathbf{AFs}_A^{D\sigma}$ the set of decisive (for the semantics σ) argumentation frameworks. The aggregation function γ is σ -**strongly decisive** if the output is always decisive :

(σ -SD) $\gamma(\mathcal{F}) \in \mathbf{AFs}_A^{D\sigma}$

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

(σ -WD) $\forall \mathcal{F} \in (\mathbf{AFs}_A^{D\sigma})^n : \gamma(\mathcal{F}) \in \mathbf{AFs}_A^{D\sigma}$

Unanimity. This specifies that if all agents agree w.r.t. some aspect of the domain (extensions, attacks, ...), for a semantics σ , then this should be reflected in the outcome.

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

(UA) $\bigcap_{k=1}^n Att(F_k) \subseteq Att(\gamma(\mathcal{F}))$

• σ -**unanimity** concerns extensions:

(σ -U) $\bigcap_{k=1}^n Ext_\sigma(F_k) \subseteq Ext_\sigma(\gamma(\mathcal{F}))$

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

(ca_σ -U) $\bigcap_{k=1}^n ca_\sigma(F_k) \subseteq ca_\sigma(\gamma(\mathcal{F}))$

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

(sa_σ -U) $\bigcap_{k=1}^n sa_\sigma(F_k) \subseteq sa_\sigma(\gamma(\mathcal{F}))$

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

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

(MAJ-A) $(|\{F_i : a \in Att(F_i)\}| > \frac{n}{2}) \Rightarrow a \in Att(\gamma(\mathcal{F}))$

• σ -**majority** concerns extensions:

(σ -MAJ) $(|\{F_i : S \in Ext_\sigma(F_i)\}| > \frac{n}{2}) \Rightarrow S \in Ext_\sigma(\gamma(\mathcal{F}))$

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

(ca_σ -MAJ) $(|\{F_i : x \in ca_\sigma(F_i)\}| > \frac{n}{2}) \Rightarrow x \in ca_\sigma(\gamma(\mathcal{F}))$

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

(sa_σ -MAJ) $(|\{F_i : x \in sa_\sigma(F_i)\}| > \frac{n}{2}) \Rightarrow x \in sa_\sigma(\gamma(\mathcal{F}))$

Closure. These properties say that the aggregation function must not invent some entity which is not in the input.

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

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

• **Attack closure** says that if one attack is in the AF in

output, this attack must be present in at least one AF in input:

(AC) $Att(\gamma(\mathcal{F})) \subseteq Att(F_1) \cup \dots \cup Att(F_n)$

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

(σ -C) $\forall S \in Ext_\sigma(\gamma(\mathcal{F})) : S \in \bigcup_{k=1}^n Ext_\sigma(F_k)$

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

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

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

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

Finally, Delobelle et al. (2015) propose a new family of properties based on the notion of Identity.

Identity. If all the AFs in the input coincide and are non-trivial, then aggregation result should be identical too.

• **Identity attack** on the attacks :

(A-ID) $Att(\gamma(F, \dots, F)) = Att(F)$

• σ -**Identity** on the extensions :

(σ -ID) $\forall F \in \mathbf{AFs}_A^{NT\sigma} : Ext_\sigma(\gamma(F, \dots, F)) = Ext_\sigma(F)$

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

(ca_σ -ID) $\forall F \in \mathbf{AFs}_A^{NT\sigma} : ca_\sigma(\gamma(F, \dots, F)) = ca_\sigma(F)$

• sa_σ -**Identity** on the sceptical inference :

(sa_σ -ID) $\forall F \in \mathbf{AFs}_A^{NT\sigma} : sa_\sigma(\gamma(F, \dots, F)) = sa_\sigma(F)$

Dunne et al. (2012) defined the properties for the case of a single AF output, whereas the merging operators may have several AFs as output. Delobelle et al. (2015) generalize the properties to ask that each AF in the output satisfies them.

Let us check what are the aggregation properties satisfied by distance-based merging operators. Since these properties do not take integrity constraints into account, we take no integrity constraint (i.e. $\mu = \top$) for the merging operators.

Proposition 8. *Given a semantics $\sigma \in \{co, pr, st, gr\}$ and two aggregation functions $\otimes = \{\Sigma, Lex\}$ and $POL = \{dg, card\}$, $\Delta_{\top}^{\otimes, d_H, \mathcal{AF}_\sigma^{POL_{avg}, \Sigma}}$ satisfies ANON, σ -U, sa_σ -U, ca_σ -C, and properties A-ID, σ -ID, ca_σ -ID and sa_σ -ID.*

• $\Delta_{\top}^{\otimes, d_H, \mathcal{AF}_{gr}^{POL_{avg}, \Sigma}}$ satisfies ca_{gr} -U, sa_{gr} -C.

• $\Delta_{\top}^{\Sigma, d_H, \mathcal{AF}_{gr}^{POL_{avg}, \Sigma}}$ satisfies gr -MAJ and ca_{gr} -MAJ.

• $\Delta_{\top}^{\Sigma, d_H, \mathcal{AF}_\sigma^{POL_{avg}, \Sigma}}$ satisfies sa_σ -MAJ.

The other properties are not satisfied.

All these results are summarized in Table 1. Basically quite few of these properties are satisfied. First, as discussed by Delobelle et al. (2015), these properties are not all that acceptable for aggregation of argumentation framework, and the existing aggregation methods do not satisfy a lot of them either. Then, as discussed in this paper, we think that the merging properties are more suited to this structured framework than these properties defined initially for an unstructured set of candidates. Basically the properties that are satisfied by our merging operators are the most important ones: anonymity, identities, unanimities, and majorities (for majoritarian aggregation functions).

Properties	Σ, dg	$\Sigma, card$	Lex, dg	$Lex, card$
ANON	✓	✓	✓	✓
σ -SNT/ σ -WNT	×	×	×	×
σ -SD/ σ -WD	×	×	×	×
UA	×	×	×	×
σ -U/ sa_σ -U	✓	✓	✓	✓
ca_σ -U	✓ ^{gr}	✓ ^{gr}	✓ ^{gr}	✓ ^{gr}
MAJ-A	×	×	×	×
σ -MAJ/ ca_σ -MAJ	✓ ^{gr}	✓ ^{gr}	×	×
sa_σ -MAJ	✓	✓	×	×
CLO/ AC / σ -C	×	×	×	×
ca_σ -C	✓	✓	✓	✓
sa_σ -C	✓ ^{gr}	✓ ^{gr}	✓ ^{gr}	✓ ^{gr}
ID	✓	✓	✓	✓

Table 1: Properties of aggregation methods. A cross \times means that the property is not satisfied, symbol \checkmark means that the property is satisfied for each $\sigma \in \{co, pr, st, gr\}$ and \checkmark^σ means that the property is satisfied only for the semantics σ .

Discussion

We want to discuss the main difference between our approach and the existing ones. The question is to know whether the main information for the agent is the set of attacks or the arguments statuses. The first option has been considered in (Coste-Marquis et al. 2007; Tohmé, Bodanza, and Simari 2008; Delobelle, Konieczny, and Vesic 2015), while we have considered the second one. We illustrate this difference on an example that is inspired by the one given by Delobelle et al. (2015). Let $\mathcal{F} = \{F_1, F_2, F_3, F_4, F_5\}$ described at Figure 5. To obtain the result of merging \mathcal{F} , an aggregation approach that relies on the attack relation leads to F_6 ; indeed, for each attack against c , the majority of the agents (at least three of them) disagree with the existence of this attack. The only attack which is accepted by the majority of the agents (in fact all of them) is the attack $b \leftarrow c$.




Figure 5: Aggregation of Attack Relations

So, with this kind of approaches, the extension of the result⁵ is $\{c, a_1, a_2, a_3\}$. This is not a problem if the attack relation is considered as a priority. But one can easily check that all the input AFs agree on the same single extension $\epsilon = \{b, a_1, a_2, a_3\}$, so postulate **(M2)** requires that this should be the extension of the resulting AF. This is exactly

⁵For all sensible semantics.

the behaviour of the merging operators which are defined in this work. At the generation step, when we consider the Hamming distance between graphs dg and the aggregation functions $\otimes = \odot = \Sigma$, the result is the set of AFs $\{F_1\}$. Formally, we have $\mathcal{AF}_\sigma^{dg}(\{\epsilon\}, \mathcal{F}) = \mathcal{AF}_\sigma^{card}(\{\epsilon\}, \mathcal{F}) = \{F_1\}$. Indeed, this set is minimal with respect to the cardinality and with respect to the sum of distances. So on this example one can check that, with our merging operators, the information on argument statuses is the most important one (and so that an extension that appear in all input AFs is selected for the result), but that we can also take into account the minimality on the attack as a second class criterion, since $F_1 = F_2$ is selected as the result since the attack $c \leftarrow a_1$ appears twice whereas all the other attacks $c \leftarrow a_i$ appear only once.

Conclusion

In this work we present the first extension-based approach to merging of AFs. Compared to previous approaches, that focused on the attack relation only, we view merging as trying to enforce minimal change at the extension level, thus treating the argument statuses as first class requirements. Inspired by propositional merging, we propose postulates **(M0)** - **(M8)** as constraints on an AF merging operator. We view postulates **(M0)** - **(M8)** as more adequate for argumentation than properties coming from Social Choice Theory, since they deal with structured pieces of information. Nonetheless, we have compared our approach to previously suggested merging operators from the literature.

The main result of the paper is a representation theorem which characterizes all operators satisfying the postulates in terms of plausibility pre-orders on candidates (sets of arguments). The concrete operators we present are based on the selection of the most plausible candidates (sets of arguments), and on the generation of the corresponding set of AFs. The generation operators allow us to ingeniously combine the minimization of the differences of the attack relations and of the cardinality of the resulting set of AFs.

An additional subtlety arises when we ask if it is possible to instantiate the merging output through a single AF: methods that work for AF revision (Diller et al. 2015) turn out to be ineffectual for merging. Thus, this is a line of research to be pursued in future work. In particular, we aim to study to which extent the constraints μ needs to be restricted to obtain operators Δ_μ that deliver a single AF. While this approach proved successful for the revision case (see (Diller et al. 2015) where an additional postulate borrowed from Horn revision (Delgrande and Peppas 2015) played a key role towards a representation theorem), recent work on Horn merging (Haret, Rümmele, and Woltran 2015) indicates that restricting the constraint in the merging operator might lead to a more substantial change in the postulates.

Acknowledgments

This work benefited from the support of the project AMANDE ANR-13-BS02-0004 of the French National Research Agency (ANR) and from the support of the Austrian Science Fund (FWF) under grants P25521 and I1102.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On The Logic Of Theory Change: Partial Meet Contraction And Revision Functions. *Journal of Symbolic Logic* 50:510–530.
- Awad, E.; Booth, R.; Tohmé, F.; and Rahwan, I. 2014. Judgment aggregation in multi-agent argumentation. *CoRR* abs/1405.6509.
- Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *Knowledge Eng. Review* 26(4):365–410.
- Baumann, R., and Brewka, G. 2015. AGM Meets Abstract Argumentation: Expansion and Revision for Dung Frameworks. In *Proc. of IJCAI'15*, 2734–2740.
- Baumann, R. 2012. What Does it Take to Enforce an Argument? Minimal Change in Abstract Argumentation. In *Proc. of ECAI'12*, 127–132.
- Bench-Capon, T. J. M., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artif. Intell.* 171(10-15):619–641.
- Bisquert, P.; Cayrol, C.; Dupin de Saint Cyr Bannay, F.; and Lagasquie-Schiex, M.-C. 2011. Change in argumentation systems: exploring the interest of removing an argument. In *Proc. of SUM'11*, 275–288. Springer-Verlag.
- Bisquert, P.; Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M. 2013. Enforcement in argumentation is a kind of update. In *Proc. of SUM'13*, 30–43.
- Boella, G.; Kaci, S.; and van der Torre, L. 2009. Dynamics in argumentation with single extensions: attack refinement and the grounded extension (extended version). In *Proc. of ArgMAS'09*, 150–159.
- Booth, R.; Kaci, S.; Rienstra, T.; and van der Torre, L. 2013. A Logical Theory about Dynamics in Abstract Argumentation. In *Proc. of SUM'13*, 148–161.
- Caminada, M., and Pigozzi, G. 2011. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems* 22(1):64–102.
- Cayrol, C.; Dupin de Saint Cyr Bannay, F.; and Lagasquie-Schiex, M.-C. 2010. Change in Abstract Argumentation Frameworks: Adding an Argument. *Journal of Artificial Intelligence Research* 38:49–84.
- Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.-C.; and Marquis, P. 2007. On the merging of Dung's argumentation systems. *Artif. Intell.* 171(10-15):730–753.
- Coste-Marquis, S.; Konieczny, S.-t.; Marquis, P.; and Ouali, M.-A. 2012a. Selecting extensions in weighted argumentation frameworks. In *Proc. of COMMA'12*, 342–349.
- Coste-Marquis, S.; Konieczny, S.-t.; Marquis, P.; and Ouali, M.-A. 2012b. Weighted attacks in argumentation frameworks. In *Proc. of KR'12*, 593–597.
- Coste-Marquis, S.; Konieczny, S.; Mailly, J.; and Marquis, P. 2014. On the Revision of Argumentation Systems: Minimal Change of Arguments Statuses. In *Proc. of KR'14*, 52–61.
- Delgrande, J. P., and Peppas, P. 2015. Belief revision in Horn theories. *Artif. Intell.* 218:1–22.
- Delobelle, J.; Konieczny, S.; and Vesic, S. 2015. On the aggregation of argumentation frameworks. In *Proc. of IJCAI'15*, 2911–2917.
- Diller, M.; Haret, A.; Linsbichler, T.; Rümmele, S.; and Woltran, S. 2015. An Extension-Based Approach to Belief Revision in Abstract Argumentation. In *Proc. of IJCAI'15*, 2926–2932.
- Doutre, S.; Herzig, A.; and Perrussel, L. 2014. A Dynamic Logic Framework for Abstract Argumentation. In *Proc. of KR'14*, 62–71.
- Dung, P. M. 1995. On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming, and n-Person Games. *Artif. Intell.* 77(2):321–357.
- Dunne, P. E.; Hunter, A.; McBurney, P.; Parsons, S.; and Wooldridge, M. 2011. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.* 175(2):457–486.
- Dunne, P. E.; Dvorák, W.; Linsbichler, T.; and Woltran, S. 2014. Characteristics of Multiple Viewpoints in Abstract Argumentation. In *Proc. of KR'14*, 72–81.
- Dunne, P. E.; Dvorák, W.; Linsbichler, T.; and Woltran, S. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.* 228:153–178.
- Dunne, P. E.; Marquis, P.; and Wooldridge, M. 2012. Argument aggregation: Basic axioms and complexity results. In *Proc. of COMMA'12*, 129–140.
- Haret, A.; Rümmele, S.; and Woltran, S. 2015. Merging in the Horn Fragment. In *Proc. of IJCAI'15*, 3041–3047.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.* 52:263–294.
- Konieczny, S., and Pérez, R. P. 2002. Merging information under constraints: a logical framework. *Journal of Logic and Computation* 12(5):773–808.
- Kontarinis, D.; Bonzon, E.; Maudet, N.; Perotti, A.; van der Torre, L.; and Villata, S. 2013. Rewriting Rules for the Computation of Goal-Oriented Changes in an Argumentation System. In *Proc. of CLIMA XIV*.
- Krümpelmann, P.; Thimm, M.; Falappa, M. A.; García, A. J.; Kern-Isberner, G.; and Simari, G. R. 2012. Selective revision by deductive argumentation. In *Proc. of TFA'11*, 147–162.
- Nouioua, F., and Würbel, E. 2014. Removed set-based revision of abstract argumentation frameworks. In *Proc. of ICTAI'14*, 784–791.
- Rahwan, I., and Simari, G. R., eds. 2009. *Argumentation in Artificial Intelligence*. Springer.
- Sakama, C. 2014. Counterfactual reasoning in argumentation frameworks. In *Proc. of COMMA'14*, 385–396.
- Tohmé, F. A.; Bodanza, G. A.; and Simari, G. R. 2008. Aggregation of attack relations: A social-choice theoretical analysis of defeasibility criteria. In *Proc. of FoIKS'08*, 8–23.