Semantic Change and Extension Enforcement in Abstract Argumentation

Sylvie Doutre[1] and Jean-Guy Mailly[2]

[1] IRIT, Université Toulouse 1 Capitole, France
doutre@irit.fr
[2] LIPADE, Université Paris Descartes, France
jean-guy.mailly@parisdescartes.fr

Abstract. Change in argumentation frameworks has been widely studied in the recent years. Most of the existing works on this topic are concerned with change of the structure of the argumentation graph (addition or removal of arguments and attacks), or change of the outcome of the framework (acceptance statuses of arguments). Change on the acceptability semantics that is used in the framework has not received much attention so far. Such a change can be motivated by different reasons, especially it is a way to change the outcome of the framework. In this paper, it is shown how semantic change can be used as a way to reach a goal about acceptance statuses in a situation of extension enforcement.

1 Introduction

Recently, the dynamics of argumentation frameworks (AFs) has received much attention [10,12,6,7,5,25,11,19,15,16,17]. Essentially, we can distinguish between two kinds of approaches for change in AFs: some of them deal with the structure of the AF (the set of arguments and the attack relation), while the other ones deal with the statuses of arguments (extensions, labellings, skeptically accepted arguments,. . .). However, a third component of the argumentation process has received almost no attention: the semantics which links the structure of the AF and the arguments statuses. Even if some approaches allow to change the semantics during the process (see for instance [6]), it is not explained *why* the semantics has to change, nor *how* the new semantics is selected. In this paper, we study these questions by focusing on extension-based semantics, that is, semantics that, when applied to an AF, produce a set of acceptable sets of arguments called extensions.

Two main reasons may motivate a change of the semantics. First, it may be required by some practical considerations. Indeed, an issue with some argumentation semantics is their high complexity. This theoretical complexity is not a practical problem if we consider some particular classes of AFs, or if the size of the AF is not too large. However, if at some point, for an agent, using some high complexity semantics is the best choice for some reason – for instance, because it guarantees the existence of at least one extension, or a number of extensions smaller than with another potential semantics –, the evolution of the AF may justify a change of the semantics. If the agent interacts with other agents in the context of a debate for instance, arguments and attacks may be added to the AF. Such additions increase the size of the AF, and they may cause the AF to leave the structural class it belongs to; this may make the computation of the extensions, and of related decision problems, not efficient anymore. A change of semantics may then be suitable.

A second reason that may motivate a change of the semantics, is as an alternative way to enforce some constraint on the acceptance statuses of arguments, or on sets of arguments. Actually, there may be limitations in given applications, which prevent to modify the attack relation and to modify the set of arguments (*e.g.* the debate the arguments and the attacks come from has ended; nothing can be added any longer). Then, if the agent has to enforce a constraint about acceptance statuses, the only component which may be modified is the semantics (that is, the way to reason about the AF). In fact, whether or not a change of the structure of the AF is possible, we show that a change of semantics can be a way to reach this goal with less change on the structure of the AF.

Main Contributions

1. We give a unified abstract framework to describe change of AFs, which encompasses all existing approaches for modifying AFs. This allows to use the same tools to analyze and extend these different approaches.
2. We extend existing work on the characteristics of extension enforcement [5], *i.e.* we provide new results about the minimal change to make on an AF to ensure that a set of arguments is (included in) an extension, w.r.t. a specific semantics.
3. We study the success rate of semantic change for extension enforcement, *i.e.* the percentage of AFs for which the result is better (w.r.t. minimal change on the AF structure) when semantic change is used. This contribution relies on the abstract framework defined in 1., and benefits from the new characteristics given in 2.

Organization of the Paper Section 2 presents background notions about abstract argumentation. Section 3 proposes a very general way to define change in argumentation frameworks, which encompasses all existing approaches. In Section 4, we show how semantic change can be used to enforce an acceptability constraint in an argumentation framework. Section 5 describes our experimental analysis of the semantic change success rate. The last section concludes the paper and describes some research tracks for future work.

2 Background Notions

[22] considers argumentation as the study of relations between arguments, without taking into account the origin of arguments or their internal structure. In this context, an argumentation framework (AF) is a directed graph $\langle A, R \rangle$ where the nodes in A are the *arguments* and the edges in R represent *attacks* between arguments. We consider only finite AFs, *i.e.* the set of arguments A is finite. $(a_i, a_j) \in R$ means that a_i attacks a_j; a_i is called an *attacker* of a_j. An argument a_i (resp. a set of arguments S) *defends* an argument a_j against its attacker a_k if a_i (resp. some argument in S) attacks a_k. The *range* of a set of arguments S w.r.t. R, denoted S_R^+, is the subset of A which contains S

and the arguments attacked by S; formally $S_R^+ = S \cup \{a_j \mid \exists a_i \in S \text{ s.t. } (a_i, a_j) \in R\}$. Different methods allow to evaluate the arguments. A common approach is to compute *extensions*, which are sets of jointly acceptable arguments. Different semantics have been defined, which yield different kinds of extensions [22,2].

Definition 1. Let $F = \langle A, R \rangle$ be an AF. A set $S \subseteq A$ is

- conflict-free w.r.t. F if $\nexists a_i, a_j \in S$ s.t. $(a_i, a_j) \in R$;
- admissible w.r.t. F if S is conflict-free and S defends each $a_i \in S$;
- a naive extension of F if S is a maximal conflict-free set (w.r.t. $\subseteq$);
- a complete extension of F if S is admissible and S contains all the arguments that it defends;
- a preferred extension of F if S is a maximal complete extension (w.r.t. $\subseteq$);
- a stable extension of F if S is conflict-free and $S_R^+ = A$;
- a grounded extension of F if S is a minimal complete extension (w.r.t. $\subseteq$);

As shortcuts, we write respectively cf, ad, na, co, pr, st, gr for these semantics. For each semantics σ, the σ-extensions of F are denoted $Ext_\sigma(F)$.

We introduce the notion of defense function[3] of a set of arguments in an AF.

Definition 2. Given an AF $F = \langle A, R \rangle$ and a set of arguments $E \subseteq A$, the *defense function of E in F* is the mapping from E and F to the set of arguments $f(E, F)$ defined by:

$$f(E, F) = \{a \in A \mid E \text{ defends } a \text{ against all its attackers}\}$$

Example 1. Let us consider the argumentation framework F_1 given at Figure 1, and let us illustrate some of the semantics.
$Ext_{ad} = \{\emptyset, \{a_1\}, \{a_4\}, \{a_4, a_6\}, \{a_1, a_3\}, \{a_1, a_4\}, \{a_1, a_4, a_6\}\}$, $Ext_{st}(F) = \{\{a_1, a_4, a_6\}\}$, $Ext_{pr}(F) = \{\{a_1, a_3\}, \{a_1, a_4, a_6\}\}$, $Ext_{co}(F) = \{\{a_1\}, \{a_1, a_3\}, \{a_1, a_4, a_6\}\}$, $Ext_{gr}(F) = \{\{a_1\}\}$.

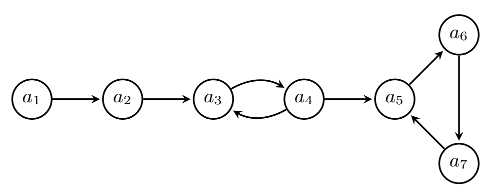

Fig. 1: The AF F_1

Let us introduce a way to measure the difference between argumentation semantics. This distance between semantics has been proposed by [20]. It relies upon the relationships between the sets of extensions that the semantics produce.

[3] This function is called *characteristic function* by [22]. We call it defense function to avoid confusion with the characteristics from [5].

Definition 3. Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a set of semantics, the *extension inclusion graph* of Σ is defined by $Inc(\Sigma) = \langle \Sigma, D \rangle$ with $D \subseteq \Sigma \times \Sigma$ such that $(\sigma_i, \sigma_j) \in D$ if and only if

- for each AF F, $Ext_{\sigma_i}(F) \subseteq Ext_{\sigma_j}(F)$;
- there is no $\sigma_k \in \Sigma$ ($k \neq i, k \neq j$) such that $Ext_{\sigma_i}(F) \subseteq Ext_{\sigma_k}(F)$ and $Ext_{\sigma_k}(F) \subseteq Ext_{\sigma_j}(F)$.

Given $\sigma_i, \sigma_j \in \Sigma$, the *$\Sigma$-inclusion difference measure* between semantics is the length of the shortest non-oriented path between σ_i and σ_j in $Inc(\Sigma)$, denoted $\delta_{Inc,\Sigma}(\sigma_i, \sigma_j)$.

Example 2. Figure 2 describes the extension inclusion graph of $\Sigma = \{cf, ad, na, st, pr, co, gr\}$. We observe, for instance, that $\delta_{Inc,\Sigma}(st, ad) = 3$, $\delta_{Inc,\Sigma}(pr, gr) = 2$, and $\delta_{Inc,\Sigma}(co, pr) = 1$.

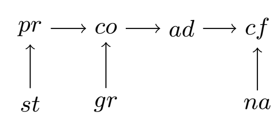

Fig. 2: Extension Inclusion Graph $Inc(\Sigma)$

3 Abstracting Change in Argumentation

We propose here an abstract definition which encompasses all methods for change in argumentation into a global family.

Definition 4. A *change operator* is a mapping χ from a multiset of AFs $\mathcal{F} = \{\!\{F_1, \ldots, F_n\}\!\}$, a formula φ from a logical language and a semantics σ, to a multiset $\mathcal{F}' = \{\!\{F'_1, \ldots, F'_k\}\!\}$ and a semantics σ'. Formally,

$$\chi(\mathcal{F}, \varphi, \sigma) = (\mathcal{F}', \sigma')$$

Most of existing operations on change in argumentation consider a single AF in the input and the output, which are obviously special cases of multiset. It is similar for approaches which consider a set as the outcome. [18] considers a profile of AFs as the input, which can be equivalently defined as a multiset since the order of the AFs in the tuple is not considered. Except [6], existing works do not consider semantic change, which means that $\sigma' = \sigma$ for these approaches. The formula represents a piece of information which is at the origin of the change (for instance in a context of belief revision [15,16] or update [25,19]). More generally, it is a constraint which has to be satisfied by the result of the operation, like an integrity constraint in a belief merging context [18]. The language of the formula is not the same depending the approach (*e.g.* each of [25,19,15,16] has its own language). Some approaches also do not use directly a formula from a logical language, but can be mapped to a formula from a given language.

For instance, adding or removing attacks and arguments [10,12] are equivalent to formulae from the language defined in [16]. Similarly, sets of arguments considered for extension enforcement [6,5,17] are special cases of the formulae defined in [15,18].

Among these approaches, some of them consider some notion of minimality, like minimal change on the attack relation [5,25,19,16], minimal change on the acceptance statuses of arguments [15,16,18], or minimal cardinality [15,18]. We can give a general definition of minimality in the change process.

Definition 5. A *minimality criterion* is a mapping from a tuple $\langle \mathcal{F}, \varphi, \sigma, \mathcal{F}', \sigma' \rangle$ to a tuple of positive real numbers $d(\langle \mathcal{F}, \varphi, \sigma, \mathcal{F}', \sigma' \rangle)$.
Given two such tuples t_1, t_2, we define $t_1 < t_2$ if the ith element of t_1 is smaller than the ith element of t_2, when i is the smallest index such that t_1 and t_2 are different.

Given a multiset of AFs $\mathcal{F} = \{\!\{F_1, \ldots, F_n\}\!\}$, a formula φ and a semantics σ, a change operator χ satisfies the minimality criterion d iff $\chi(\mathcal{F}, \varphi, \sigma) = (\mathcal{F}', \sigma')$ and $d(\langle \mathcal{F}, \varphi, \sigma, \mathcal{F}', \sigma' \rangle)$ is minimal.

Obviously, the simplest minimality criteria can be defined with a single number, so $d(\langle \mathcal{F}, \varphi, \sigma, \mathcal{F}', \sigma' \rangle)$ is a tuple of length 1. For instance, we instantiate this definition with extension enforcement operators [6,5,17].

Definition 6. Given an AF $F = \langle A, R \rangle$ and a set of arguments $E \subseteq A$, a *strict* (resp. *non-strict*) *enforcement operator* is a change operator which maps $\mathcal{F} = \{\!\{F\}\!\}$, a formula $\varphi_E = \bigwedge_{a_i \in E} a_i$ and a semantics σ to $\mathcal{F}' = \{\!\{F'\}\!\}$ and σ' such that $E \in Ext_{\sigma'}(F')$ (resp. $\exists \epsilon \in Ext_{\sigma'}(F')$ with $E \subseteq \epsilon$).

An enforcement is *minimal* iff if satisfies the minimality criterion

$$d(\langle \mathcal{F}, \varphi, \sigma, \mathcal{F}', \sigma' \rangle) = \langle d_H(\mathcal{F}, \mathcal{F}') \rangle$$

where d_H is the Hamming distance between graphs [4] [5].

We say that F' is an enforcement of E in F. We use $\varphi_E = \bigwedge_{a_i \in E} a_i$ to specify that the set E is the enforcement request; this is reminiscent of the logical encodings used in [17,26].

Some change operators use more complex minimality criteria, which combine m simple criteria. In this case, we can represent it with a m-length tuple; this is the case of *e.g.* [15,16,18].

4 Extension Enforcement and Semantic Change

In this section, we study how semantic change can be useful for extension enforcement. We first recall the definition of the five existing enforcement approaches. Then we show

[4] The Hamming distance between two graphs $F_1 = \langle A_1, R_1 \rangle$ and $F_2 = \langle A_2, R_2 \rangle$ is the cardinality of the symmetric difference between R_1 and R_2; in other words, in the present case, it is the number of attacks that it is necessary to add/remove from one graph to get the other.

[5] Since here $\mathcal{F}, \mathcal{F}'$ are singletons, the Hamming distance between graphs can be directly used. For other kinds of change operators, it should be generalized to multisets.

on intuitive examples that changing the semantics can permit to enforce an extension with fewer change on the structure (or even *without* any structural change). Finally, we extend Baumann's study on minimal change depending on the semantics, and we define a more general class of enforcement operators which reach our goal: perform extension enforcement with minimal structural change *by semantic change*.

4.1 Extension Enforcement Operators

In the first work on extension enforcement [6], it is considered that everything which appears in the current AF cannot be changed. The authorized changes are the addition of arguments, and possibly of attacks concerning at least one new argument. This kind of change is called a *normal expansion*. Special cases of normal expansion are called *strong expansion* and *weak expansion*. A strong expansion (resp. weak expansion) is an expansion which adds only strong arguments (resp. weak arguments), which are arguments that cannot be attacked by (resp. cannot attack) the previous arguments.

Definition 7. Let F, F' be two AFs such that F' is a strict (resp. non-strict) enforcement of a set of arguments E in F.

- If F' is a normal expansion of F, then the change from F to F' is a *strict* (resp. *non-strict*) *normal enforcement*.
- If F' is a strong expansion of F, then the change from F to F' is a *strict* (resp. *non-strict*) *strong enforcement*.
- If F' is a weak expansion of F, then the change from F to F' is a *strict* (resp. *non-strict*) *weak enforcement*.

Then, [17] considers new approaches which, on the opposite, question the attack relation between existing arguments. Two operators are proposed.

Definition 8. Let $F = \langle A, R \rangle, F' = \langle A', R' \rangle$ be two AFs such that F' is a strict (resp. non-strict) enforcement of the set of arguments E in F.

- If $A = A'$ and $R \neq R'$, then the change from F to F' is a *strict* (resp. *non-strict*) *argument-fixed enforcement*.
- If $A \subseteq A'$, then the change from F to F' is a *strict* (resp. *non-strict*) *general enforcement*.

In all these approaches, it is considered that

- either the semantics does not change in the enforcement;
- or the new semantics is given as a parameter of the operator: it is not specified why the semantics should change, nor why this particular semantics should be the new one.

We use $Nor_x, Str_x, Weak_x, Fix_x$ and Gen_x to denote these enforcement methods, with $x \in \{s, ns\}$ corresponding to strict and non-strict.

4.2 Minimal Structural Change through Semantic Change

Example 3. Let us consider again the AF F_1 given at Figure 1. We want to enforce the set $E = \{a_1, a_3\}$ as an extension. We consider that the agent is currently using the stable semantics. Obviously, structural change is required if the agent does not change the semantics. But we have seen previously that E is already an extension of F if we consider, for instance, the preferred or the complete semantics. So if the agent considers a change of semantics, the enforcement can be realized without any change on the structure.

Of course, in some situations, only switching the semantics may not be sufficient to reach the goal, if none of the possible semantics leads to build extensions which are consistent with this goal. In this case, and even if structural change is permitted, then the semantic change can still be a means to minimize the structural change required to reach the goal. Indeed, even if structural changes are permitted (or required), it can be costly for the agents to perform such changes. Such modifications of the set of arguments and of the set of attacks may then have to be limited.

The minimal change problem for extension enforcement has already been studied in [5], for a subset of the possible enforcement approaches. First, it only considers some particular target semantics (stable, preferred, complete, admissible). Also, the argument-fixed enforcement operators is not considered. Finally, only non-strict enforcement is characterized. For each pair of these semantics and enforcement operators, the minimal number of changes (addition or removal of attacks) to reach an enforcement is called the *characteristic*. This characteristic is a natural number when the enforcement is possible; $+\infty$ means that the enforcement is impossible under the given semantics.

We continue this study of characteristics and we give here some results for argument-fixed enforcement. We first need to introduce some notations.

Definition 9. Given an AF $F = \langle A, R \rangle$, and $X \subseteq A$,

- $R_\downarrow(F, X) = R \cap (X \times X)$ for any $X \subseteq A$;
- $na(F, X) = \{a_i \in A \setminus X \mid \forall a_j \in X, (a_i, a_j) \notin R$ and $(a_j, a_j) \notin R\}$
- $ad(F, X) = \{a_i \in A \setminus X \mid \exists a_j \in X, (a_i, a_j) \in R$ and $\forall a_j \in X, (a_j, a_i) \notin R\}$
- $st(F, X) = \{a_i \in A \setminus X \mid \forall a_j \in X, (a_j, a_i) \notin R\}$.

Proposition 1. *Let $F = \langle A, R \rangle$ be an AF, and $E \subseteq A$. The characteristic of strict argument-fixed enforcement for $\sigma \in \{cf, ad, st, co, pr, na\}$ is defined by the function $V^F_{\sigma,Fix_s}(E)$:*

$V^F_{cf,Fix_s}(E) = |R_\downarrow(F, E)|$
$V^F_{na,Fix_s}(E) = |R_\downarrow(F, E)| + |na(F, E)|$
$V^F_{ad,Fix_s}(E) = |R_\downarrow(F, E)| + |ad(F, E)|$
$V^F_{st,Fix_s}(E) = |R_\downarrow(F, E)| + |st(F, E)|$
$V^F_{co,Fix_s}(E) = \min\{|R' \Delta R| + |R_\downarrow(F', E)| \mid f(E, F') = E, F' = \langle A, R' \rangle\}$
$V^F_{pr,Fix_s}(E) = \min\{|R' \Delta R| + |R_\downarrow(F', E)| \mid E \subseteq f(E, F'), \forall E \subset E' \subseteq A, E' \not\subseteq f(E', F'), F' = \langle A, R' \rangle\}$

We observe that these results are in line with the complexity results from [26]. Indeed, these characteristics suggest polynomial-time algorithm to compute the minimal enforcement of E under cf, na, ad and st semantics. Obtaining a better formulation for the other characteristics is still challenging.

Proposition 2. *Let $F = \langle A, R\rangle$ be an AF, and $E \subseteq A$. The characteristic of non-strict argument-fixed enforcement for $\sigma \in \{cf, ad, st, co, pr, na\}$ is defined by the function $V^F_{\sigma,Fix_{ns}}(E)$:*
$V^F_{na,Fix_{ns}}(E) = V^F_{cf,Fix_{ns}}(E) = |R_{\downarrow}(F, E)|$
$V^F_{ad,Fix_{ns}}(E) = \min(\{|R_{\downarrow}(F, E')| + |ad(F, E')| \mid E \subseteq E' \subseteq A\})$
$V^F_{st,Fix_{ns}}(E) = \min(\{|R_{\downarrow}(F, E')| + |st(F, E')| \mid E \subseteq E' \subseteq A\})$
$V^F_{pr,Fix_{ns}}(E) = V^F_{co,Fix_{ns}}(E) = V^F_{ad,Fix_{ns}}(E)$

We notice that these results are reminiscent of the characteristics for general enforcement [5].

Observation 1. For $Op \in \{Nor, Str, Weak\}$, the characteristic is trivial for conflict-free and naive semantics: either the set E is conflict-free, then the characteristic is 0; or E is not conflict-free, then the characteristic is $+\infty$.

Now, we generalize the definition of enforcement operators to take into account semantic change.

Definition 10. Let $F = \langle A, R\rangle$ be an AF, σ a semantics, Σ be a set of semantics, and $E \subseteq A$. Given $Op \in \{Nor, Str, Weak, Fix, Gen\}$ and $x = s$ (resp. $x = ns$), the minimal change enforcement of E in F w.r.t. Op_x is defined as $\chi(\{\!\{F\}\!\}, \varphi_E, \sigma) = (\{\!\{F'\}\!\}, \sigma')$ with $\sigma' \in \Sigma$, such that $E \in Ext_{\sigma'}(F')$ (resp. $\exists \epsilon \in Ext_{\sigma'}(F')$ s.t. $E \subseteq \epsilon$), and the criterion $\langle V^F_{\sigma,Op_x}(E), \delta_{Inc,\Sigma}(\sigma, \sigma')\rangle$ is satisfied.

This means that contrary to previous works on extension enforcement, the target semantics is not a parameter of the enforcement operator. It is chosen to guarantee that:

- the characteristic (*i.e.* the structural change) is minimal;
- in the case when several semantics have the same characteristic, the chosen one should minimize the semantic change.

Example 4. Let us come back to the AF F_1 described at Figure 1. We want to enforce the set $E = \{a_1, a_3\}$ as an extension, with $\sigma = st$ the semantics currently used by the agent. E is not a stable extension, neither the grounded extension or a naive extension. However, it is a preferred, complete, admissible and conflict-free extension. This means that

- for every $\sigma' \in \{pr, co, ad, cf\}$, $V^F_{\sigma,Op_x}(E) = 0$ for every Op_x;
- for every $\sigma' \in \{st, gr, na\}$, $V^F_{\sigma,Op_x}(E) > 0$ for every Op_x.

This guarantees that the result of the enforcement (whatever the operator Op_x) is the AF F_1 itself, with one of the semantics $\{pr, co, ad, cf\}$. We observe that $\delta_{Inc,\Sigma}(st, pr) = 1$, $\delta_{Inc,\Sigma}(st, co) = 2$, $\delta_{Inc,\Sigma}(st, ad) = 3$ and $\delta_{Inc,\Sigma}(st, cf) = 4$, so the new semantics is the preferred semantics. Formally, the result of enforcing E in F_1 is

$$Op_x(\{\!\{F_1\}\!\}, \bigwedge_{a_i \in E} a_i, st) = (\{\!\{F_1\}\!\}, pr)$$

We use here $\delta_{Inc,\Sigma}$ to illustrate our approach, but other difference measures between semantics could be used to define minimal semantic change. The inclusion graph that we use here is a particular case of relation graph as defined in [20]. Some other interesting notions of relation graphs could be used to define distances between semantics, like intertranslatability graphs [23] or skepticism relations [3]. [20] also mentions other approaches, based on the properties satisfied by the semantics, or based on the actual set of extensions of an AF w.r.t. the different semantics. This offers a wide range of possibilities to define minimal semantic change.

Observation 2. Our approach cannot give a worse result, w.r.t. structural change, than the classical enforcement approaches (by "classical", we mean approaches without semantic change, or with a given target semantics). Moreover, we can identify some basic cases for which our approach is sure to give a better result than classical approaches. For instance, as illustrated by Example 4, when the set E to be enforced is not a σ-extension of the considered AF F (with σ the current semantics), but E is known to be a σ'-extension of F, with σ' one of the possible alternative semantics. In this situation, it is guaranteed that enforcing E in F with our semantic change-based approach is possible without any structural change, while classical approaches do not permit this.

5 Empirical Study

In this section, we present an empirical study of the success of semantic change for extension enforcement. We have computed the result of some enforcement requests for a large set of AFs (using the strict argument-fixed enforcement approach), w.r.t. different semantics ($\Sigma = \{ad, st, co, na\}$), and for each pair $(\sigma_1, \sigma_2) \in \Sigma \times \Sigma$, we have compared $V^F_{\sigma_1, Fix_s}$ and $V^F_{\sigma_2, Fix_s}$. When $V^F_{\sigma_1, Fix_s}$ is significantly higher than $V^F_{\sigma_2, Fix_s}$ for a given AF F, this means that semantic change is relevant for this AF, w.r.t. this pair of semantics and enforcement operator. Indeed, in this case, changing the semantics from σ_1 to σ_2 allows to reach one's goal (enforcing a set of arguments E) with a lower cost (w.r.t. change of the graph). In the following subsections, we first present in detail our experimental protocol, then we provide an analysis of our results.

5.1 Protocol

We have used the AFs and enforcement requests from [26], which are available online. They provide AFs with different size of arguments $|A| \in \{50, 100, 150, 200, 250, 300\}$. The AFs are generated following the Erdös-Rényi model [24]. For $p \in \{0.05, 0.1, 0.2, 0.3\}$, each pair $(a_i, a_j) \in A \times A$ has a probability p to belong to the attack relation R. For each $|A|$ and each p, five AFs have been generated. Finally, for each AF, five sets of arguments $E \subset A$ have been randomly generated for each $|E|/|A| \in \{0.05, 0.1, 0.2, 0.3\}$. This means that for each $|A|$, 400 enforcement problem instances $(F = \langle A, R \rangle, E \subset A)$ have been generated.

For all these enforcement requests, we have computed the result of the argument-fixed strict enforcement for $\sigma \in \{na, ad, stb, co\}$. Enforcement under the naive semantics has been done through a software that we have developped in Java. For the other semantics, we have used Pakota, the enforcement solver provided by [26].[6]

The experiments have been done on a 64bits Ubuntu 16.04 system, equipped with 8Gio of RAM and a CPU Intel Core i5 with 3.20GHz. The time limit was set to 10 minutes.

5.2 Analysis of the Results

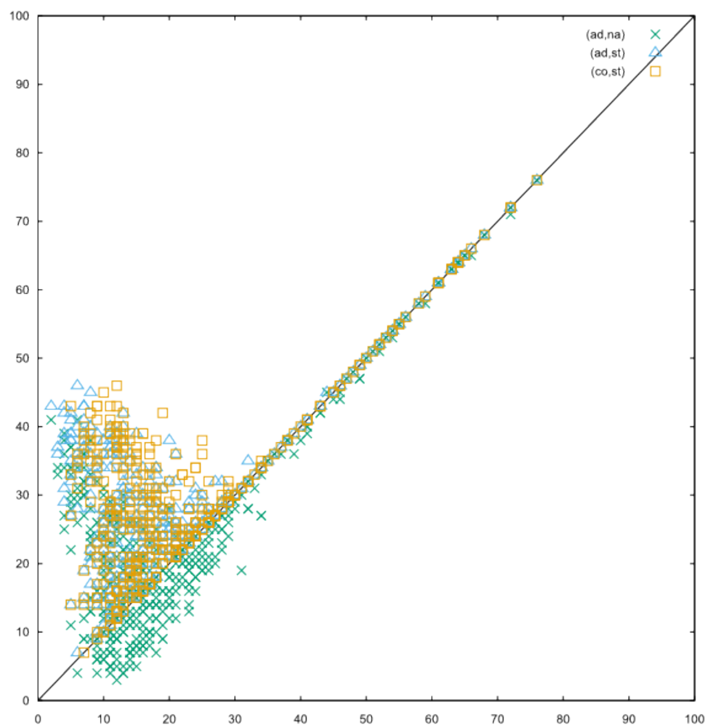

Fig. 3: Comparing Minimal Change Depending on the Semantics, for AFs with 50 Arguments

Figure 3 presents our results for a subset of the instances, namely the AFs with $|A| = 50$ and the associated enforcement requests $E \subset A$. We only present the results for this class of AFs for a matter of readability. Indeed, for the other values of $|A|$, the results appear to be remarkably similar. Also, we only present 3 of the 6 possible

[6] Pakota also provides the possibility to execute enforcement under the preferred semantics. Because of the higher complexity of the enforcement problem under the preferred semantics, our experiment has encountered a high number of timeouts. For this reason, we exclude preferred semantics of our empirical analysis for now.

combinations of semantics: (ad, st) (represented by $\triangle$), (ad, na) (represented by $\times$) and (co, st) (represented by $\square$). For each of these combinations (σ_1, σ_2), each point represents an instance (*i.e.* a pair $(F = \langle A, R \rangle, E \subseteq A)$), such that the point abscissa is the minimal change to enforce E in F w.r.t. σ_1, and its ordinate is the value for the enforcement w.r.t. σ_2. So, a point situated under the diagonal represents an instance for which the minimal change to perform the enforcement w.r.t. σ_1 is higher than the minimal change to perform the enforcement w.r.t. σ_2 (and vice-versa for the points above the diagonal). We observe that semantic change actually brings something to extension enforcement. Indeed for most of the instances, the points are situated far from the diagonal, which means that they can benefit from semantic change. On the opposite, the points situated on the diagonal represent instances for which semantic change does not improve the ”quality” of enforcement.

Let us mention the fact that we have similar results for the pairs of semantics (st, na) and (co, na). Only the pair (ad, co) results in points close to the diagonal for a high proportion of the instances. For $|A| \in \{100, 150, 200, 250, 300\}$, we observe similar results. Let us still mention that the higher the value of $|A|$, the higher the proportion of instances with a ratio close to 1. But even for $|A| = 300$, there is still a significant amount of instances which benefit from semantic change (*i.e.* instance with a significative difference between $V^F_{\sigma_1, Fix_s}$ and $V^F_{\sigma_2, Fix_s}$).[7] Figure 4 presents, for each $|A|$ and each pair of semantics, the percentage of instances for which the ratio $V^F_{\sigma_1, Fix_s}/V^F_{\sigma_2, Fix_s}$ is smaller than 0.9 or greater than 1.1, *i.e.* the percentage of instance for which semantic change is successful.

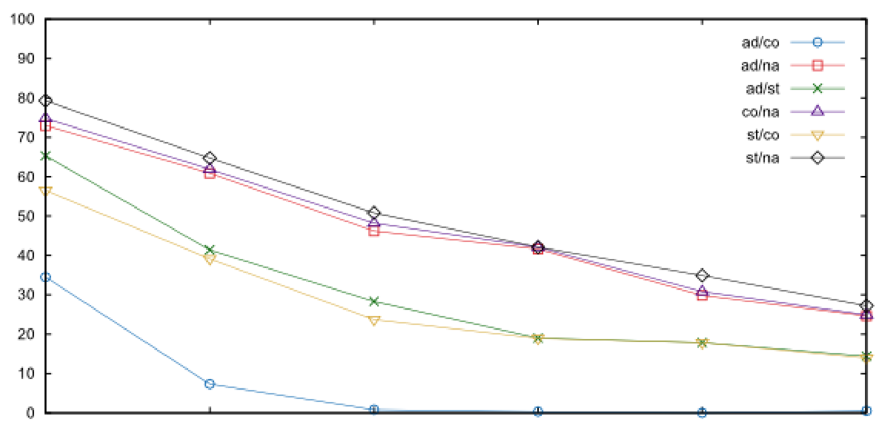

Fig. 4: Success Percentage for Different Semantic Change Situations

[7] A complete description and analysis of our experiments, including the instances, the enforcement system, and the curves for every value of $|A|$ and every pair (σ_1, σ_2) is available online: `http://www.math-info.univ-paris5.fr/~jmailly/expSemChange`

6 Conclusion

This paper addresses particular aspects of the *dynamics of argumentation frameworks*. Most of the existing approaches in this domain concern either a change of the structure of an AF, or a change on the acceptance statuses of arguments (both being related). We argue that it makes sense in some applications to permit the agent to change her reasoning process, which is represented by the acceptance semantics. This change can be motivated by a need of computational efficiency (requirement of a lower complexity), or by properties to be enforced on the set of extensions (e.g. requirement of some arguments to be accepted), with a minimal change of the graph structure.

Such a change in the reasoning process is related to what is discussed in [8,9]. Roughly speaking, the idea is that an agent can be able to use different reasoning processes, such as one which is harder to compute and probably more rational, and another one which is easier to compute and based on some less rational concepts (for instance, there can be some bias due to the agent's perception of the source of information). Semantic change in argumentation can be conducted by similar ideas.

In this paper, we have first defined a very abstract framework to describe change in argumentation. This framework is useful to describe and analyze the different approaches for argumentation dynamics with the same tools. Then we have instantiated this framework for a specific (and well-studied) family of change operators for AFs: extension enforcement. We show that allowing an agent to change the semantics when performing an extension enforcement is useful in some situations, since this semantic change cannot provide a worse result (w.r.t. the number of modifications of the graph) than "classical" enforcement, and can even provide better results. This claim is grounded on the new study of characteristics. We have conducted an experimental study which shows the impact of semantic change on a large set of instances.

Several interesting questions have arisen from this work. Naturally, we want to complete our study of characteristics and our experiments with more semantics. The ideal semantics [21], the prudent semantics [14] or the SCC-recursive semantics [4] are good candidates. Determining the missing characteristics (for instance, the characteristics of the strict versions of operators studied by Baumann in [5]) is also an important future work. Since the difference between semantics is here evaluated in the setting of the well-known extension-based semantics, the extension of our approach to labelling-based semantics seems to be quite immediate. On the contrary, semantic change for ranking-based semantics [1] requires a deeper investigation. Regarding our experimental study, we want to explore more in depth the impact of the different parameters on the semantic change, for instance the size of the AF, the size of the set of arguments to be enforced, and the probability of attacks. We have considered here the Erdös-Rényi model, which captures an interesting graph structure, and which has already been the object of other studies [26]. We plan to conduct similar studies with other families of graphs [13] to determine whether the impact of semantic change is different for these families. Also, we want to extend extension enforcement systems to benefit from the study of characteristics: computing the characteristics for a list of enforcement operators and a list of semantics, we can choose the best operator and semantics to enforce a set with minimal change of the graph.

Finally, we want to study the impact of semantic change on some operations which return a set [15,18]. In these papers, the outcome of the operation represents some uncertain result (intuitively, the set is interpreted as a “disjunction” of AFs). Our goal is to determine whether semantic change can help to reduce the cardinality of the set (i.e. reduce the uncertainty of the result).

References

1. Amgoud, L., Ben-Naim, J.: Ranking-based semantics for argumentation frameworks. In: Proc. SUM’13. pp. 134–147 (2013)
2. Baroni, P., Caminada, M., Giacomin, M.: An introduction to argumentation semantics. Knowledge Eng. Review (2011)
3. Baroni, P., Giacomin, M.: Skepticism relations for comparing argumentation semantics. Int. J. Approx. Reasoning 50(6), 854–866 (2009)
4. Baroni, P., Giacomin, M., Guida, G.: SCC-recursiveness: a general schema for argumentation semantics. Artif. Intell. (2005)
5. Baumann, R.: What does it take to enforce an argument? minimal change in abstract argumentation. In: Proc. ECAI’12. pp. 127–132 (2012)
6. Baumann, R., Brewka, G.: Expanding argumentation frameworks: Enforcing and monotonicity results. In: Proc. COMMA’10. pp. 75–86 (2010)
7. Bisquert, P., Cayrol, C., Dupin de Saint-Cyr, F., Lagasquie-Schiex, M.C.: Change in argumentation systems: exploring the interest of removing an argument. In: Proc. SUM 2011. pp. 275–288 (2011)
8. Bisquert, P., Croitoru, M., Dupin de Saint-Cyr, F.: Four ways to evaluate arguments according to agent engagement. In: Proc. BIH’15. pp. 445–456 (2015)
9. Bisquert, P., Croitoru, M., Dupin de Saint-Cyr, F.: Towards a dual process cognitive model for argument evaluation. In: Proc. SUM’15. pp. 298–313 (2015)
10. Boella, G., Kaci, S., van der Torre, L.: Dynamics in argumentation with single extensions: Attack refinement and the grounded extension. In: Proc. AAMAS’09. pp. 1213–1214 (2009)
11. Booth, R., Kaci, S., Rienstra, T., van der Torre, L.: A logical theory about dynamics in abstract argumentation. In: Proc. SUM 2013. pp. 148–161 (2013)
12. Cayrol, C., Dupin de Saint-Cyr, F., Lagasquie-Schiex, M.C.: Change in abstract argumentation frameworks: Adding an argument. J. Artif. Intell. Res. 38, 49–84 (mai 2010)
13. Cerutti, F., Giacomin, M., Vallati, M.: Generating structured argumentation frameworks: AFBenchGen2. In: Proc. COMMA’16. pp. 467–468 (2016)
14. Coste-Marquis, S., Devred, C., Marquis, P.: Prudent semantics for argumentation frameworks. In: Proc. ICTAI 2005. pp. 568–572 (2005)
15. Coste-Marquis, S., Konieczny, S., Mailly, J.G., Marquis, P.: On the revision of argumentation systems: Minimal change of arguments statuses. In: Proc. KR’14 (2014)
16. Coste-Marquis, S., Konieczny, S., Mailly, J.G., Marquis, P.: A translation-based approach for revision of argumentation frameworks. In: Proc. JELIA’14 (2014)
17. Coste-Marquis, S., Konieczny, S., Mailly, J.G., Marquis, P.: Extension enforcement in abstract argumentation as an optimization problem. In: Proc. IJCAI’15 (2015)
18. Delobelle, J., Haret, A., Konieczny, S., Mailly, J.G., Rossit, J., Woltran, S.: Merging of abstract argumentation frameworks. In: Proc. KR’16. pp. 33–42 (2016)
19. Doutre, S., Herzig, A., Perrussel, L.: A dynamic logic framework for abstract argumentation. In: Proc. KR’14. pp. 62–71 (2014)
20. Doutre, S., Mailly, J.G.: Quantifying the difference between argumentation semantics. In: Proc. COMMA’16 (2016)

21. Dung, P.M., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. Artif. Intell. 171(10-15), 642–674 (2007)
22. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. Artif. Intell. 77(2), 321–357 (1995)
23. Dvorák, W., Spanring, C.: Comparing the expressiveness of argumentation semantics. In: Proc. COMMA'12. pp. 261–272 (2012)
24. Erdös, P., Rényi, A.: On random graphs i. Publicationes Mathematicae pp. 290–297 (1959)
25. de Saint-Cyr, F.D., Bisquert, P., Cayrol, C., Lagasquie-Schiex, M.: Argumentation update in YALLA (yet another logic language for argumentation). Int. J. Approx. Reasoning 75, 57–92 (2016)
26. Wallner, J.P., Niskanen, A., Järvisalo, M.: Complexity results and algorithms for extension enforcement in abstract argumentation. In: Proc. AAAI'16. pp. 1088–1094 (2016)