

On the Revision of Argumentation Systems: Minimal Change of Arguments Statuses

Sylvie Coste-Marquis Sébastien Konieczny Jean-Guy Mailly Pierre Marquis

CRIL, CNRS – Université d’Artois, Lens, France
{coste,konieczny,maily,marquis}@cril.fr

Abstract

In this paper, we investigate the revision of argumentation systems à la Dung. We focus on revision as minimal change of the arguments status. Contrarily to most of the previous works on the topic, the addition of new arguments is not allowed in the revision process, so that the revised system has to be obtained by modifying the attack relation only. We introduce a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system. We show how AGM belief revision postulates can be translated to the case of argumentation systems. We provide a corresponding representation theorem in terms of minimal change of the arguments statuses. Several distance-based revision operators satisfying the postulates are also pointed out, along with some methods to build revised argumentation systems. We also discuss some computational aspects of those methods.

Introduction

In this paper, we investigate the revision issue for abstract argumentation systems à la Dung (Dung 1995). Argumentation systems are directed graphs, where nodes correspond to arguments and arcs to attacks between arguments. In such systems, the status (acceptance) of each argument depends on the chosen acceptability semantics (grounded, preferred, stable – among others).

In his book, Gärdenfors (1988) introduced abstractly belief change as the operation allowing to change the epistemic status of a piece of information with respect to the epistemic state of an agent. There are three possible statuses: accepted, rejected or undetermined. And revision, contraction and expansion are defined as the possible transitions between these statuses, as illustrated by Figure 1.

In order to instantiate Gärdenfors’ general definition of belief change within Dung’s argumentation theory, it is first necessary to define what are the available pieces of information and what these statuses mean. In Dung’s argumentation theory, the basic pieces of information are the arguments of the system, and their statuses depend on the acceptability semantics under consideration. Thus, following Gärdenfors, it




Figure 1: Gärdenfors’ epistemic transitions

does not make sense to study the revision of argumentation systems directly on the attack graph, independently of any semantics. Stated otherwise, the revision of a given argumentation system under two different semantics may easily lead to two different results. For instance, in the case of the argumentation system AF given in Figure 2, we note that under the stable and preferred semantics, d belongs to every extension, whereas d does not belong to any extension in the case of the grounded semantics. So, revising AF in order to accept d does not need any change for the stable or the preferred semantics, but a change is required for the grounded semantics.




Figure 2: Revising AF may lead to different results.

In this paper, we focus on revision as minimal change of the arguments statuses. To be more precise, under a chosen semantics, given an argumentation system and a revision formula expressing how the statuses of some arguments has to be changed, we want to derive argumentation systems which satisfy the revision formula, and are such that the corresponding extensions are as close as possible to the extensions of the input system. In most of the previous works on the topic, the change in the input system is a

modification of the set of arguments. Contrarily to those works, we do not allow the addition of new arguments in the revision process, so that the revised system has to be obtained by modifying the attack relation, only. Especially, the revision formula does not indicate *why* the statuses of arguments have changed, which makes sense in cases when the agent learns an unjustified information from a trustworthy source. More details can be found in Section “Discussion”.

Minimal change of the attack graph can be considered as a criterion for defining the revised systems, but the acceptance statuses of arguments is more fundamental information. Accordingly, ensuring a minimal change of these statuses is more important than (and not compatible with)¹ ensuring a minimal change of the attack graph.

To ensure such a minimal change of arguments statuses, we define revision as two step process: first, we revise extensions, without considering the attack relation. Then, from the output of this first step, we generate argumentation systems such that their extensions are the expected ones. This second step may use minimal change on the argumentation graph, as a second criterion.

The rest of the paper is organized as follows. After a short introduction to Dung’s theory of abstract argumentation, we introduce a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system. Then we show how AGM belief revision postulates can be translated to the case of argumentation systems. We provide a corresponding representation theorem in terms of minimal change of the arguments statuses. Several distance-based revision operators satisfying the postulates are pointed out. Then we present methods for associating argumentation systems with the obtained sets of extensions. We discuss some computational aspects of the revision process. Proofs are omitted for space reasons.

Preliminaries

We start with a very short introduction to Dung’s theory of argumentation (see (Dung 1995) for more details). A (finite) argumentation system (also referred to as an argumentation framework) is a pair $AF = \langle A, R \rangle$ where A is a (finite and non empty) set of so-called arguments and R is a binary relation over A (a subset of $A \times A$). In the following, A is supposed to contain at least two elements and the attack relation R is supposed to be irreflexive, i.e., self-contradicting arguments are rejected. \mathbf{AFS}_A denotes the set of all such systems on the set of arguments A .

Arguments and attacks are considered in an abstract way as atomic concepts. Especially we do not assume any underlying logical setting in which arguments and attacks would be defined.

An argument $a \in A$ is acceptable with respect to a set of arguments $S \subseteq A$ whenever it is defended by S , i.e., for every $b \in A$ s.t. $(b, a) \in R$, there exists $c \in S$ such that $(c, b) \in R$. We say that a subset S of A is conflict-free if and

only if for every $a, b \in S$, we have $(a, b) \notin R$. A subset S of A is admissible for AF if and only if S is conflict-free and acceptable with respect to S . “Solutions” of an argumentation systems are sets S of arguments accepted together. Some semantics σ (especially, the complete, the preferred, the stable, and the grounded semantics) can be considered for capturing formally this notion, and each of them gives rise to a specific notion of extension. For instance:

- S is a complete extension of AF if and only if it is an admissible set and every argument which is acceptable with respect to S belongs to S ,
- S is a preferred extension of AF if and only if it is maximal (with respect to set inclusion) in the set of admissible sets for AF ,
- S is a stable extension of AF if and only if S is conflict-free and $\forall a \in A \setminus S, \exists b \in S$ such that $(b, a) \in R$,
- S is the grounded extension of AF if and only if it is the smallest element (with respect to set inclusion) among the complete extensions.

$Ext_\sigma(AF)$ denotes the set of extensions of AF for the semantics σ . In this work, we focus on the skeptical policy to define the epistemic status of an argument. Thus, the epistemic status of any argument $a \in A$ with respect to the epistemic state represented by AF is given by: a is accepted if a belongs to every $\varepsilon \in Ext_\sigma(AF)$, a is refused if a does not belong to any $\varepsilon \in Ext_\sigma(AF)$, and a is undefined in the remaining case.

Moreover, we use the following notation about minimal elements of a set. For any pre-order \leq over a set E , $<$ denotes the strict part of \leq and \simeq denotes the indifference relation associated with \leq . Given a set E and a pre-order \leq , the minimal elements of E w.r.t. \leq are $\min(E, \leq) = \{e \in E \mid \nexists e' \in E, e' < e\}$.

On Revision Formulae

We want to define a revision setting for Dung’s argumentation systems in which sophisticated revision formulae can be taken into account, and not only the fact that a given argument should be accepted or refused. To this end, we consider a logical language \mathcal{L}_A , where negation is used to denote the fact that a given argument should be refused, and formulae can be connected using conjunction and disjunction.

Definition 1. Given $A = \{\alpha_1, \dots, \alpha_k\}$ a set of arguments, \mathcal{L}_A is the language generated by the following context-free grammar in BNF:

$$\begin{aligned} arg & ::= \alpha_1 \mid \dots \mid \alpha_k \\ \Phi & ::= arg \mid \neg \Phi \mid (\Phi \wedge \Phi) \mid (\Phi \vee \Phi) \end{aligned}$$

For instance, $\varphi_1 = (a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)$ expresses that in the revised epistemic state, a must be accepted and b and c must be both accepted or both refused. The epistemic status of such a formula φ from \mathcal{L}_A in an argumentation system $AF \in \mathbf{AFS}_A$ for a given semantics σ is given by:

Definition 2. Let $\varepsilon \subseteq A$ and $\varphi \in \mathcal{L}_A$. The concept of satisfaction of φ by ε , noted $\varepsilon \models \varphi$, is defined inductively as follows:

¹See Section “Revision at the System Level” for details.

- If $\varphi = a \in A$, then $\varepsilon \vdash \varphi$ iff $a \in \varepsilon$,
- If $\varphi = (\varphi_1 \wedge \varphi_2)$, $\varepsilon \vdash \varphi$ iff $\varepsilon \vdash \varphi_1$ and $\varepsilon \vdash \varphi_2$,
- If $\varphi = (\varphi_1 \vee \varphi_2)$, $\varepsilon \vdash \varphi$ iff $\varepsilon \vdash \varphi_1$ or $\varepsilon \vdash \varphi_2$,
- If $\varphi = \neg\psi$, $\varepsilon \vdash \varphi$ iff $\varepsilon \not\vdash \psi$.

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

- φ is accepted w.r.t. AF, noted $AF \vdash_\sigma \varphi$, if $\varepsilon \vdash \varphi$ for every $\varepsilon \in \text{Ext}_\sigma(AF)$,
- φ is refused w.r.t. AF, noted $AF \vdash_\sigma \neg\varphi$, if $\varepsilon \vdash \varphi$ for no $\varepsilon \in \text{Ext}_\sigma(AF)$,
- φ is undefined w.r.t. AF in the remaining case.

Inference \vdash_σ can be extended to the case of a set S of argumentation systems by considering $\text{Ext}_\sigma(S) = \bigcup_{AF \in S} \text{Ext}_\sigma(AF)$.

This language, used in next sections to define the revision of an argumentation framework, allows to change the status of an information φ to accepted (revise by φ) or rejected (revise by $\neg\varphi$), but not to undetermined. It is normal: Gärdenfors defines the change of status of a belief from accepted or rejected to undetermined as a *contraction*, as explained in the Introduction (see Fig. 1).

From now on we call *candidate*² any subset ε of A . Candidates can be interpreted as *interpretations* of revision formulae. Continuing the previous example, if $A = \{a, b, c\}$, φ_1 is satisfied by the candidates from $\{\{a\}, \{a, b, c\}\}$. Thus, for the grounded semantics, φ_1 is accepted w.r.t. AF_1 with $R_1 = \{(b, c), (c, b)\}$ but is refused w.r.t. AF_2 with $R_2 = \{(a, b), (b, a)\}$.

We define consistency in a classical way:

Definition 3. Given a formula φ , \mathcal{A}_φ denotes the set of candidates satisfying φ . φ is said to be consistent iff $\mathcal{A}_\varphi \neq \emptyset$.

In the general case, \mathcal{A}_φ is not the set of all σ -extensions of an AF in \mathbf{AFs}_A . Consider for instance, $A = \{a, b, c\}$ and $\varphi_1 = (a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)$. $\mathcal{A}_{\varphi_1} = \{\{a\}, \{a, b, c\}\}$, and there is no AF in \mathbf{AFs}_A such that $\text{Ext}_\sigma(AF) = \mathcal{A}_{\varphi_1}$ for $\sigma = \text{grounded}$, $\sigma = \text{preferred}$ or $\sigma = \text{stable}$.

In this case, it is enough to choose two argumentation frameworks to cover the extensions $\{\{a\}, \{a, b, c\}\}$ (for instance, in the first AF a attacks b and c , and in the second AF the attack relation is empty). Note that in the general case, increasing the number of frameworks is not enough to capture the expected extensions. In order to characterize formulae that can be associated to a set of frameworks and a semantics, a concept of σ -representability can be defined as follows:

Definition 4. A set C of candidates is σ -representable iff there exists a set S of argumentation systems in \mathbf{AFs}_A such that $C = \text{Ext}_\sigma(S)$.

A formula $\varphi \in \mathcal{L}_A$ is σ -representable iff \mathcal{A}_φ is σ -representable.

When $A = \{a, b, c\}$, $\varphi_1 = (a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)$ is σ -representable for $\sigma = \text{grounded}$, $\sigma = \text{preferred}$ or $\sigma = \text{stable}$ since $\{\{a\}, \{a, b, c\}\} = \text{Ext}_\sigma(AF_3) \cup \text{Ext}_\sigma(AF_4)$

²That is to say “candidate to be an extension”.

where $R_3 = \{(a, b), (a, c)\}$ and $R_4 = \emptyset$. Contrastingly, $\varphi_2 = \neg a \wedge \neg b \wedge \neg c$ is grounded-representable and preferred-representable but not stable-representable. $\varphi_3 = a \wedge \neg a$ is neither grounded-representable nor preferred-representable, but is stable-representable (consider AF_5 such that $R_5 = \{(a, b), (b, c), (c, a)\}$).

A form of consistency can be defined to take account for the semantics:

Definition 5. Given σ a semantics, a formula $\varphi \in \mathcal{L}_A$ is σ -consistent iff φ is consistent and φ is σ -representable.

From σ -consistency, we define a notion of model:

Definition 6. Given a formula $\varphi \in \mathcal{L}_A$ and a semantics σ , the set of models of φ is defined by

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

A last point about formulae is the definition of equivalence. Two formulae $\varphi, \psi \in \mathcal{L}_A$ are said σ -equivalent, noted $\varphi \equiv_\sigma \psi$, if and only if $\mathcal{A}_\varphi^\sigma = \mathcal{A}_\psi^\sigma$.

Revision at the Extension Level

In order to define revision operators, our approach follows a two-step process. Intuitively, the process first selects from models of φ those as close as possible to the σ -extensions of AF . This selection has to ensure the minimal change of arguments statuses. Then, the second step generates the argumentation systems such that the union of their σ -extensions precisely coincides with the selected candidates.

We define a revision operator on argumentation systems as a mapping associating a set of argumentation systems with the input argumentation system and the input revision formula:

Definition 7. Given any set of arguments A , a revision operator on argumentation systems \star is a mapping from $\mathbf{AFs}_A \times \mathcal{L}_A$ to $2^{\mathbf{AFs}_A}$.

Clearly, the result of the revision of an argumentation system is not a unique argumentation system in the general case, but a set of argumentation systems. The reason is quite simple: there can be several possible results which have exactly the same maximum plausibility. So in this case there is no reason to select just one of them (we will return to this point later on). If this is problematic for a particular application, a selection function can be used as a tie-break rule for ensuring the unicity of the result (just like, for instance, the maxichoice selection function considered in AGM belief revision (Gärdenfors 1988)).

Of course, each mapping from $\mathbf{AFs}_A \times \mathcal{L}_A$ to $2^{\mathbf{AFs}_A}$ is not a reasonable revision operator. For instance, the constant, yet trivial operator defined by $AF \star \varphi = \emptyset$ should be discarded.

In order to identify interesting revision operators, we have to identify the logical properties which guarantee a rational behaviour. Such an axiomatic approach is standard in

logic, and the AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985; Katsuno and Mendelzon 1991) have been pointed out for characterizing valuable revision operators in a logical setting. As in (Qi, Liu, and Bell 2006), we can revisit these postulates in a set-theoretic framework, here suited to the argumentation case.

Let \mathcal{S} be a set of argumentation systems AF in \mathbf{AFs}_A . The counterpart of AGM postulates in the argumentation case is given by:

(AE1) $Ext_\sigma(AF \star \varphi) \subseteq \mathcal{A}_\varphi^\sigma$

(AE2) If $Ext_\sigma(AF) \cap \mathcal{A}_\varphi^\sigma \neq \emptyset$,
then $Ext_\sigma(AF \star \varphi) = Ext_\sigma(AF) \cap \mathcal{A}_\varphi^\sigma$

(AE3) If φ is σ -consistent, then $Ext_\sigma(AF \star \varphi) \neq \emptyset$

(AE4) If $\varphi \equiv_\sigma \psi$, then $Ext_\sigma(AF \star \varphi) = Ext_\sigma(AF \star \psi)$

(AE5) $Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi^\sigma \subseteq Ext_\sigma(AF \star (\varphi \wedge \psi))$

(AE6) If $Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi^\sigma \neq \emptyset$,
then $Ext_\sigma(AF \star (\varphi \wedge \psi)) \subseteq Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi^\sigma$

(AE1) states that the σ -extensions of the resulting set of argumentation systems must be among the models of φ . (AE2) demands that if there are σ -extensions of the input system satisfying φ , then the resulting σ -extensions must coincide with them. (AE3) requires the resulting set of σ -extensions to be non-empty as soon as φ is σ -consistent. (AE4) concerns the irrelevance of syntax: the revision by two formulae must be identical if the formulae are equivalent. The last two postulates (AE5) and (AE6) express a minimal change principle with respect to the arguments statuses: changes of the statuses of the arguments are expected to be minimal with respect to the input system. In particular, these postulates give the expected behaviour of the operator when an argumentation system is revised by a conjunction of formulae.

Interestingly, as in the logical case, we can derive a representation theorem which characterizes exactly the revision operators satisfying the postulates in a more constructive way. To this end, we first need to present a counter-part of the notion of faithful assignment (Katsuno and Mendelzon 1991) in the argumentation setting:

Definition 8. A faithful assignment is a mapping associating any argumentation system $AF = \langle A, R \rangle$ (under a semantics σ) with a total pre-order \leq_{AF}^σ on the set of candidates such that:

1. if $\varepsilon_1 \in Ext_\sigma(AF)$ and $\varepsilon_2 \in Ext_\sigma(AF)$,
then $\varepsilon_1 \simeq_{AF}^\sigma \varepsilon_2$,
2. if $\varepsilon_1 \in Ext_\sigma(AF)$ and $\varepsilon_2 \notin Ext_\sigma(AF)$,
then $\varepsilon_1 <_{AF}^\sigma \varepsilon_2$.

The representation theorem can then be stated as follows:

Proposition 1. Given a semantics σ , a revision operator \star satisfies the rationality postulates (AE1) - (AE6) iff there exists a faithful assignment which maps every system $AF = \langle A, R \rangle$ to a total pre-order \leq_{AF}^σ so that

$$Ext_\sigma(AF \star \varphi) = \min(\mathcal{A}_\varphi^\sigma, \leq_{AF}^\sigma).$$

This theorem is important for defining operators satisfying the rationality postulates, as the ones presented in the next section.

Distance-Based Revision

Let us now present some (pseudo-)distance-based revision operators satisfying the rationality postulates (AE1) - (AE6).

Let d be any pseudo-distance³ on 2^A , for instance, the Hamming distance given by $d_H(\varepsilon_1, \varepsilon_2) = |(\varepsilon_1 \setminus \varepsilon_2) \cup (\varepsilon_2 \setminus \varepsilon_1)|$. Given $\varepsilon \in 2^A$ and $\mathcal{E} \subseteq 2^A$, d can be extended to a “distance” between ε and \mathcal{E} , by stating that $d(\varepsilon, \mathcal{E}) = \min_{\varepsilon' \in \mathcal{E}} d(\varepsilon, \varepsilon')$. For any argumentation system $AF \in \mathbf{AFs}_A$, this distance induces a total pre-order between candidates $\varepsilon_1, \varepsilon_2 \in 2^A$ given by

$$\varepsilon_1 \leq_{AF}^{\sigma, d} \varepsilon_2 \text{ iff } d(\varepsilon_1, Ext_\sigma(AF)) \leq d(\varepsilon_2, Ext_\sigma(AF)).$$

On this ground, revision operators can be defined by:

Definition 9. Let σ be any given semantics. A pseudo-distance-based revision operator \star^d is any revision operator for which there exists a pseudo-distance d on 2^A such that for every AF and every φ , we have $Ext_\sigma(AF \star^d \varphi) = \min(\mathcal{A}_\varphi^\sigma, \leq_{AF}^{\sigma, d})$.

Proposition 2. Let σ be any semantics. Any pseudo-distance-based revision operator \star^d satisfies the rationality postulates (AE1) - (AE6).

Let us now define another family of pseudo-distance-based operators, which take advantage of labellings. Let us first recall that, instead of using extensions, the solutions of an argumentation system can be expressed using the concept of labelling (Caminada 2006).

Formally, a labelling is a mapping L associating a label *in*, *undec* or *out* with every argument of the set A . The stronger notion of reinstatement labelling depends on the attack relation R : an argument a is labelled *in* iff every argument attacking a is *out*; an argument a is *out* iff there exists an argument *in* attacking a ; an argument is *undec* iff it is neither *in* nor *out*. These reinstatement labellings correspond to Dung’s complete extensions in a bijective way. Thus, for any complete extension ε , the associated reinstatement labelling is such that every argument $a \in \varepsilon$ is *in*, every argument attacked by an argument *in* is *out*, every other argument is *undec*. Conversely, for every reinstatement labelling L , the corresponding extension $E(L)$ is the set of arguments labelled *in* by the labelling. All usual semantics can also be encoded with labellings.

We introduce some notation: L_φ is the set of labellings L such that $E(L) \in \mathcal{A}_\varphi^\sigma$. Given a set of labellings Lab , $E(Lab) = \{E(L) | L \in Lab\}$. Finally, given a system AF and a semantics σ , $Labs_\sigma(AF)$ denotes the set of labellings corresponding to the σ -extensions of AF .

Labellings, which bring richer information than extensions, can be used to define interesting pseudo-distance-based revision operators. Indeed, consider the following notion of edition pseudo-distance:

³A pseudo-distance d on a set S is defined as a binary relation on S which satisfies:

- $d(x, y) = d(y, x)$;
- $d(x, y) = 0$ iff $x = y$.

Definition 10. Let m, n, o be three integers and let L_1 and L_2 be two labellings.

An edition pseudo-distance $d_{(m,n,o)}$ between labellings is defined as:

$$d_{(m,n,o)}(L_1, L_2) = \sum_{a \in A} ad(L_1(a), L_2(a)),$$

where

- $ad(in, in) = ad(out, out) = ad(undec, undec) = 0$
- $ad(in, out) = ad(out, in) = m$
- $ad(in, undec) = ad(undec, in) = n$
- $ad(out, undec) = ad(undec, out) = o$

Proposition 3. Let m, n, o be three integers. $d_{(m,n,o)}$ is a pseudo-distance.

Interestingly, these edition pseudo-distances are not necessarily neutral or symmetric. We call neutral an edition pseudo-distance such that $ad(in, undec) + ad(undec, out) = ad(in, out)$ and symmetric a pseudo-distance such that $ad(in, undec) = ad(undec, out)$. Defining non-symmetric edition pseudo-distances is a way for instance to favor acceptance of arguments over rejection (see Example 1).

For any pseudo-distance $d_{\mathcal{L}}$ between labellings, we can define a pre-order $\leq_{AF}^{\sigma, d_{\mathcal{L}}}$ between labellings as we did it for candidates:

$$L_1 \leq_{AF}^{\sigma, d_{\mathcal{L}}} L_2 \text{ iff } d_{\mathcal{L}}(L_1, L_{\sigma}(AF)) \leq d_{\mathcal{L}}(L_2, L_{\sigma}(AF)).$$

Definition 11. Let σ be any given semantics. A labelling-pseudo-distance-based revision operator $\star_{d_{\mathcal{L}}}$ is any revision operator for which there exists a pseudo-distance $d_{\mathcal{L}} = d_{(m,n,o)}$ on 2^A such that for every AF and every φ , we have $Labs_{\sigma}(AF \star_{d_{\mathcal{L}}} \varphi) = \min(L_{\varphi}, \leq_{AF}^{\sigma, d_{\mathcal{L}}})$.

The following example illustrates the impact of the chosen pseudo-distance on the revised system:

Example 1. Let σ be the stable semantics. We revise the system AF_6 below by the formula $\varphi = (\neg d \wedge \neg e)$.




Figure 3: The system AF_6

$Ext_{\sigma}(AF_6) = \{\{a, d, e\}, \{b, d, e\}\}$, the associated stable labellings are $\{(a, in), (b, out), (c, out), (d, in), (e, in)\}$ and $\{(a, out), (b, in), (c, out), (d, in), (e, in)\}$. When we revise AF_6 by φ using the pseudo-distance-based operator induced by the pseudo-distance $d_{(1,9,10)}$ on labellings, the obtained result is a system with the following labellings: $\{(a, in), (b, out), (c, out), (d, out), (e, out)\}$ and $\{(a, out), (b, in), (c, out), (d, out), (e, out)\}$. When the pseudo-distance $d_{(9,1,10)}$ is used, we get $\{(a, in), (b, out), (c, out), (d, undec), (e, undec)\}$ and $\{(a, out), (b, in), (c, out), (d, undec), (e, undec)\}$ as labellings of the result systems.

If the second step of the process, i.e. the generation of the resulting argumentation systems, as expected, takes account for the labellings, the structure of the resulting graphs will be different: when the refused arguments are out, it means that there exists an attack from an accepted argument to a refused argument. When the arguments are undec, those attacks do not exist.

With the first pseudo-distance $d_{(1,9,10)}$, it is cheaper to change an argument from in to out than to undec. Such a pseudo-distance allows for choosing candidates which refuse arguments. Contrarily, the pseudo-distance $d_{(9,1,10)}$ allows for choosing candidates which accept more arguments.

Like operators based on extensions, pseudo-distance-based operators using labelling exhibit good logical properties:

Proposition 4. Let σ be any semantics. Any labelling-pseudo-distance-based revision operator $\star_{d_{\mathcal{L}}}$ satisfies the rationality postulates (AE1) - (AE6).

Revision at the System Level

The operators defined in the previous section focus on the candidates that are as close as possible to the extensions of the input system. However, they do not indicate how to generate the corresponding argumentation systems, i.e., the argumentation systems such that the union of their extensions coincides with the selected candidates. This task is the second step in the definition of the revision operator.

In order to achieve this task, we consider a mapping \mathcal{AF}_{σ} from 2^{2^A} to $2\mathcal{AFs}_A$, called generation operator, that associates with any set \mathcal{C} of candidates a set of argumentation systems such that $Ext_{\sigma}(\mathcal{AF}_{\sigma}(\mathcal{C})) = \mathcal{C}$.

An important point we would like to discuss is the fact that a revision operator \star outputs a set of argumentation systems, and not a single argumentation system in the general case. Actually, this is a consequence of the expressiveness of the language of revision formulae we want to consider. In order to illustrate it, consider $A = \{a, b, c, d\}$, and AF_7 as represented in Figure 4.




Figure 4: The system AF_7




Figure 5: Revision of AF_7

The extensions of AF_7 are the same ones for the stable and preferred semantics, $Ext_\sigma(AF_7) = \{\{a, b\}, \{a, c\}\}$. Let $\varphi_4 = (-b \vee c) \wedge (-c \vee b)$. Observe that b and c play symmetric roles, both in AF_7 and in φ_4 . When computing the result of the revision with the revision operator based on Hamming distance between candidates, we obtain two candidates $\{a\}$ and $\{a, b, c\}$. We present in the following generation operators leading to the two corresponding argumentation systems AF_8 (corresponding to candidate $\{a\}$) and AF_9 (corresponding to candidate $\{a, b, c\}$). Clearly, choosing one of these systems would require to accept some arbitrariness given the symmetric roles of b and c .

Let us now show that for every semantics σ , generation operators \mathcal{AF}_σ do exist.

Proposition 5. *Whatever the semantics σ , for every non-empty set \mathcal{C} of candidates from 2^A , such that $\emptyset \notin \mathcal{C}$, there exists a finite set $\mathcal{S} \subseteq \mathbf{AFs}_A$ such that $\mathcal{C} = Ext_\sigma(\mathcal{S})$.*

For instance, for every candidate C , we can build the naive system $\langle A, R_C \rangle$ with $R_C = \{(x, y) \in A \times A \mid x \in C, y \notin C\}$.

So now we can define our revision operators:

Definition 12. *Given a semantics σ , a faithful assignment that matches every argumentation system to a total pre-order \leq_{AF}^σ , and a generation operator \mathcal{AF}_σ , the corresponding revision operator \star is defined by:*

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

One of the key results of the paper is that:

Proposition 6. *Every revision operator \star defined following Definition 12 satisfies the postulates (AE1)-(AE6).*

By construction, our revision operators are ensured to deal with minimality of change of arguments statuses, but not with minimality of change of the attack relation. Indeed, the rationality postulates ask for preserving as much as possible the statuses of arguments in the input system: doing so while ensuring that the revision formula is satisfied does not usually imply a minimal change of the attack relation, and vice-versa. As a matter of illustration, consider the argumentation systems AF_{10} , AF_{11} , AF_{12} , and AF_{13} .




Figure 6: Minimal Change

Suppose that our goal is to reject e , that is to get a system so that e does not appear in any extension. So we consider the revision formula $\varphi_5 = \neg e$. A minimal change on the attack relation of AF_{10} leads to AF_{11} or AF_{12} : each of them differs with AF_{10} on a single attack. This contrasts with AF_{13} since the change on the attack relation required to go from AF_{10} to AF_{13} is strictly greater than the change on the attack relation required to go from AF_{10} to AF_{12} . Each of these four systems has a unique extension for the usual semantics⁴: $\{b, d, e\}$ for AF_{10} , $\{b, c, d\}$ for AF_{11} , and $\{b, d\}$ for AF_{12} and AF_{13} . Hence, the change on the statuses of arguments achieved when going from AF_{10} to AF_{12} or AF_{13} is strictly smaller than the change on the statuses of arguments achieved when going from AF_{10} to AF_{11} .

During the generation process, minimization can actually be considered in at least two ways: either minimizing change on the attack relation, or minimizing the number of output systems. In fact, these ways can be combined, either with a more important role to minimal change of the attack relation, or with a more important consideration for minimization of the cardinality.

Thus, a notion of minimal change on the attack relation can be defined through a notion of pseudo-distance dg on the attack relation. Such a pseudo-distance can be for instance the Hamming distance, given by $dg_H(AF_1, AF_2) = |(\mathcal{R}_1 \setminus \mathcal{R}_2) \cup (\mathcal{R}_2 \setminus \mathcal{R}_1)|$. The dg_H distance between two argumentation systems corresponds to the number of attacks that must be added or removed to make them identical. But we can consider more elaborated edition pseudo-distances such as those given in (Coste-Marquis et al. 2007). Each pseudo-distance dg induces a pre-order between argumentation systems, defined by $AF_1 \leq_{AF}^{dg} AF_2$ iff $dg(AF_1, AF) \leq dg(AF_2, AF)$. We can easily extend this notion to a distance between a system AF and a set of systems AFs : $dg(AF, AFs) = \min_{AF_i \in AFs} (dg(AF, AF_i))$.

To give priority to minimal change on the attack relation, we define a generation operator that builds sets of argumentation systems which cover the candidates; then chooses the ones which minimize a function of the pseudo-distance dg ; and finally retains the sets which are minimal in terms of cardinality.

Definition 13. *Given \mathcal{C} a set of candidates, σ a semantics, dg a pseudo-distance between graphs and AF an argumentation system, $\mathcal{AF}_\sigma^{dg, AF}$ is defined as:*

$$\mathcal{AF}_\sigma^{dg, AF}(\mathcal{C}) = \bigcup \{AFs \in sets \mid card(AFs) \text{ is minimal}\}$$

with

$$sets = \{AFs \mid Ext_\sigma(AFs) = \mathcal{C} \text{ and } \sum_{AF_i \in AFs} dg(AF, AF_i) \text{ is minimal}\}.$$

⁴For instance, for the complete, the preferred, the stable and the grounded semantics.

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 14. Given \mathcal{C} a set of candidates, σ a semantics, dg a pseudo-distance between graphs and AF an argumentation system, $\mathcal{AF}_\sigma^{card,AF}$ is defined by:

$$\mathcal{AF}_\sigma^{card,AF}(\mathcal{C}) = \bigcup \{AFs \in sets \mid \sum_{AF_i \in AFs} dg(AF, AF_i) \text{ is minimal}\}$$

with

$$sets = \{AFs \mid Ext_\sigma(AFs) = \mathcal{C} \text{ and } card(AFs) \text{ is minimal}\}.$$

These two approaches are not equivalent:

Proposition 7. The generation operators $\mathcal{AF}_\sigma^{dg,AF}$ and $\mathcal{AF}_\sigma^{card,AF}$ are distinct.

Example 2. Let us now give an example of revision with the approaches previously defined. The input system AF_{14} is given on Figure 7.



Figure 7: The framework AF_{14}

Its unique stable extension is $\{a, b, c\}$. The revision formula is $\varphi_6 = (a \vee b) \wedge (\neg a \vee \neg b)$, the revision operators are \star_{dg} such that

$$AF \star_{dg} \varphi = \mathcal{AF}_\sigma^{dg,AF}(\min(\mathcal{A}_\varphi^\sigma, \leq_{AF}^{\sigma, d_H}))$$

and \star_{card} such that

$$AF \star_{card} \varphi = \mathcal{AF}_\sigma^{card,AF}(\min(\mathcal{A}_\varphi^\sigma, \leq_{AF}^{\sigma, d_H}))$$

both based on the Hamming distance on candidates and the Hamming distance on attack relations. Each one uses one of the previously defined generation operators.

Let us first compute the revised candidates. It is easy to find that $\{a, c\}$ and $\{b, c\}$ are the minimal models of φ_6 with respect to the Hamming distance and the stable extension of the input system.

Now we present the result for the two revision operators. When minimizing the change on attack relation, the generation step produces two argumentation systems, AF_{15} and AF_{16} , each one with a single difference from the input graph.



Figure 8: $AF_{14} \star_{dg} \varphi_6$

Contrastingly, the revision of AF_{14} with operator \star_{card} gives as output a unique argumentation system AF_{17} with two differences with respect to the Hamming distance on the attack relation.




Figure 9: $AF_{14} \star_{card} \varphi_6$

While the two approaches exemplified here use the sum to aggregate the pseudo-distances, any aggregation function can be used instead. For instance, the min, the max, or any OWA (Ordered Weighted Average (Yager 1988)). These kind of aggregation function allows to combined distance and cardinality without giving priority to one of them. For instance, a specific ordered weighted average OWA_b is given by:

- $v(S) = (dg(AF_1, AF), \dots, dg(AF_k, AF))^5$ with $S = \{AF_1, \dots, AF_n\}$ and k the cardinal of the largest set (the vectors corresponding to smaller sets are normalized by adding the appropriate number of zeroes in front of the vector).

- $w_i = 2^{i-1}$

- $OWA_b(E) = \sum_{i=1}^k w_i v(E)[i]$

With the OWA_b function, a set of frameworks E_1 such that the vector of pseudo-distances is $v_1 = (1, 1, 4)$ is less preferred than a set E_2 with $v_2 = (1, 2, 3)$, because $OWA_b(E_1) = 19 > OWA_b(E_2) = 17$.

A set E_2 , with three argumentation systems, can also be preferred to a set E_3 with only two systems, if those two systems are too far from the input systems. For instance, if the vector of pseudo-distances is $(1, 4)$, once normalized to $v_3 = (0, 1, 4)$, $OWA_b(E_3) = 18$, and so E_2 is still preferred in spite of its larger cardinality.

It is worthwhile to note that aggregation functions can also be used alone to define a generation operator. Given an aggregation function f , we define a pre-order \leq_f such that $E_1 \leq_f E_2$ iff $f(E_1) \leq f(E_2)$. For any aggregation function on sets of argumentation systems, a generation operator is AF_σ^f given by

$$\mathcal{AF}_\sigma^f(\mathcal{C}) \in \min(\{E = \{AF_i\} \mid Ext_\sigma(E) = \mathcal{C}\}, \leq_f).$$

All proposed approaches cannot guarantee to produce a single set of argumentation systems. For instance, two sets may have the same cardinality and the same pseudo-distance from the input system with respect to the pseudo-distance dg .

The result can then be defined following one of the two options below:

- The result is defined as the union of all the sets. The reason is that these sets represent the uncertain result of the revision, so we keep all of them to avoid arbitrary choice. The main default of this method is that the size of the result is increased.
- A tie-break rule is used to select a single set of AFs. The agent is obliged to do an arbitrary choice. Note that it

⁵We suppose that the vector is sorted in increasing order.

does not prevent the revision operator from satisfying rationality postulates (because the postulates deal with extensions, not with attack relations).

Some Computational Aspects

A first interesting question concerns the size of the output of revision operators. The first step of the revision can lead to an exponential number of candidates, in terms of the number n of arguments in the input system. It is directly related to the revision formula. Given a pseudo-distance-based revision operator, the size of the output depends also on the generation operator which is used at the second step of the process. In the worst case, the number of argumentation systems which are generated is exponential in n :

Proposition 8. *If \star is a revision operator based on the generation operator $\mathcal{AF}_\sigma^{dg, AF}$, then the size of $AF \star \varphi$ can be exponential in $|A|$.*

The complexity of the inference problem has also to be identified. Given a revision operator \star for argumentation systems and a semantics σ , the *inference problem* from a revised argumentation system is the following decision problem:

- **Input:** An argumentation system AF on A , and two formulae $\varphi, \psi \in \mathcal{L}_A$.
- **Query:** Does $AF \star \varphi \sim_\sigma \psi$ hold?

Unsurprisingly, provided that σ ensures the existence of an extension for every AF , the inference problem from a revised argumentation system $AF \star \varphi$ is at least as hard as the inference problem from AF . In formal terms:

Proposition 9. *Let \mathbf{C} be a complexity class which is closed under polynomial-time reductions. Suppose that \star satisfies (AE1) to (AE6), and that the semantics σ ensures the existence of an extension for every AF . If the inference problem from an argumentation system is \mathbf{C} -hard, then the inference problem from a revised argumentation system is \mathbf{C} -hard as well.*

Clearly enough, it can be the case that the inference problem from a revised argumentation system $AF \star \varphi$ is strictly harder than the inference problem from AF (unless $\mathbf{P} = \mathbf{NP}$). For instance, under the restriction when the queries ψ are restricted to arguments (or more generally, CNF formulae on A), it is easy to show that the inference problem from AF w.r.t. the grounded semantics can be solved in polynomial time. Contrastingly:

Proposition 10. *Suppose that \star satisfies (AE1) and (AE3). The inference problem from a revised argumentation system w.r.t. the grounded semantics is \mathbf{coNP} -hard, even under the restriction when the queries ψ are restricted to CNF formulae on A .*

Our results show that the revision of argumentation systems is comparable to the revision of propositional formulae from a computational point of view. Especially, it may lead to harder computational problems: on the one hand, the revision of an argumentation system may require exponentially many systems for being represented (this is reminis-

cent to the non-compilability of some belief revision operators (Cadoli et al. 1999)); on the other hand, inference may also become harder (Nebel 1998).

Related Work

Some previous works have already considered the change issue for argumentation systems à la Dung. Thus, Boella, Kaci and Van der Torre, (2009a; 2009b) have studied abstraction and refinement principles. An abstraction is a reduction of the attack relation or of the set of arguments, whereas a refinement is the addition of attacks or arguments to the system. The authors focused on the study of semantics which ensure the existence of a unique extension (for instance, the grounded extension), and they formulated some principles of the form “if we do this particular change, then the extension of the result is like this”. They identified some principles satisfied by the grounded semantics.

Cayrol, Dupin de Saint-Cyr and Lagasquie-Schiex (2010) studied the addition of an argument to an argumentation system. They stated some properties that can be satisfied when a change occurs in an argumentation system, and pointed out those which are satisfied (and under which conditions) when an argument (and the attacks concerning it) is added to the graph. With Bisquert, they did a similar study about the deletion of an argument (Bisquert et al. 2011).

In (Kontarinis et al. 2013), the authors studied the problem of minimal change to satisfy a goal based on arguments acceptance. The main difference between this work and our approach is that minimal change for them refers to the number of changes to perform on the attack relation, while we focus on arguments statuses.

Baumann (2012) also studied the minimal change problem in abstract argumentation. He reported some bounds on the number of modifications of the attack relation to make so as to enforce a given set of arguments. These bounds depend on the semantics and the type of change allowed.

Discussion

An abstract argumentation system à la Dung is defined by a set of arguments and an attack relation. All the works discussed in the previous section generate changes in an argumentation system by adding new arguments, while trying to minimize the modifications of the attack relation. Contrastingly, if we consider that the meaning of an argumentation system is given by the set of arguments that are accepted or rejected, then minimal change means to minimize the change on the (status of the) set of arguments.

Giving the priority on the minimization of the attack relation or on the minimization of the arguments status is really a question of considering the status of the arguments as first-class citizen of argumentation systems or only as a by-product of the graph. Whereas this last option received considerable attention these last years, there is no work except this one which concentrates on the first optic.

In addition, previous works also suppose that one can add as many arguments as one wants in order to modify the status of some arguments. In some cases this is perfectly sensible, but in other cases it is difficult to assume that such

arguments are available. Consider for instance the case of big society debates, where political parties, economists, and other specialists have already put forward all the arguments in favor of or against some decision (for example whether the state has to increase or decrease individual taxation). If a political leader wants to change the current decision, then he will need to be very brilliant in order to find an argument that has not been already pointed out by experts. More probably she will rather try to change the beliefs (or preferences) of the people on the fact that some arguments do or do not attack other ones.

Let us now mention some contexts where allowing only changes of the attack relation is very natural.

A first example of application for the revision process without adding new arguments is the reception of a unjustified but trustworthy information (which is a particular case of argument from authority). This scenario is frequent in applications of argumentation on social network debates (Gabbriellini and Torroni 2013). When an agent A initiates a debate about an argument α , if another agent B does not agree with A about α but considers that A is trustworthy, B has to revise her beliefs to accept α . In this case, agent B can change her beliefs even if agent A has not introduced a new argument in the debate. So, B has to reconsider the attacks between some arguments, but not the set of arguments itself.

A second example concerns applications of argumentation on public opinion, and is related to the society debates motivation we discussed above: suppose that an argumentation system represents the opinion of some groups of agents, where an attack between arguments exists if the majority of the group supports it. If a group leader wants to modify the statuses of arguments, then she can perform a revision of the input system. The resulting argumentation systems may help her to determine the attacks she has to focus on so as to modify the majoritarian opinion.

A third context concerns preference-based argumentation (see (Amgoud and Cayrol 2002)). In such argumentation systems some arguments attacks each other (in particular if the arguments are based on logical formulae and rebuttal, attack is symmetric), and the preference relation determines if an attack succeeds or not. So it is possible to modify the attack relation just by modifying the preferences of the agent. A similar case of revision can occur with value-based argumentation frameworks (Bench-Capon 2002): each argument is mapped to a value, and a value can be “stronger” than another. Comparison of values can lead an attack to fail. In this case, a change of values leads to a change of the (succeeding) attacks.

A final point we want to discuss is the fact that our revision methods output a set of argumentation systems instead of a single one. At the beginning of Section “Revision at the System Level”, we advocated the fact that this is necessary to avoid arbitrary choices. We want to add that obtaining a set as result of a revision process is just usual in most belief revision settings. It is important to note that the canonical representation of AGM contraction/revision operators by use of relational partial-meet functions (Alchourrón, Gärdenfors, and Makinson 1985) defines the result of the process as a set of minimal theories. It turns out

that the language used makes it possible to produce a single theory from this set using intersection (conjunction). But in languages where this is not possible it seems natural to keep a set as result. For instance (Fagin et al. 1986) defines *flocks* that are the set of logical databases which result of the revision of a single logical database. Flocks have also been used as sets of possible results for combination/merging operators (Baral, Kraus, and Minker 1991; Baral et al. 1992; Konieczny 2000).

Conclusion and Future Work

In this paper, we investigated the revision problem for abstract argumentation systems à la Dung. We focused on revision as minimal change of the arguments statuses. We introduced a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system. We showed how AGM belief revision postulates can be translated to the case of argumentation systems. We provided a corresponding representation theorem in terms of minimal change of the arguments status, and pointed out several pseudo-distance-based revision operators satisfying the postulates. We investigated some computational aspects of revision of argumentation systems.

We are currently encoding our revision operators by representing argumentation systems with logical constraints (in a similar way to (Besnard and Doutre 2004)), so as to be able to benefit from the power of constraint solvers to compute revised systems. At the time of writing this paper, the revision of the stable extensions of an AF by a logical constraint is encoded. Some future work is to define an encoding for other semantics, and to encode some generation operators. Informations about this work is available here: <http://www.cril.fr/DynArgs/revision.html>.

There are some things still to do concerning this topic. Let us quote some of them now.

Designing revision approaches which allow to combine the two kinds of revision operators of argumentation systems would be useful: the ones that allow only addition of new arguments and no change of the attack relation, and the ones that allow only changes of the attack relation and no addition of new arguments. Allowing both changes in a sensible way is an interesting question, especially in the presence of constraints (stating for instance that some arguments must remain accepted, or that some attacks cannot be changed).

Associating a minimal set of argumentation systems with a set of candidates is another important issue, not only for our revision purpose. It is related to the problem of realizability (Dunne et al. 2013), where the question is to find a (unique) argumentation system that corresponds to a set of candidates. This problem can also be studied in the case of labellings, and used for the generation of argumentation systems from a set of labellings, exploiting labelling-distance-based revision operators defined in this paper.

Finally, it is interesting to study if some restriction on the faithful assignment could allow to guarantee that the output of the revision operator is a single AF, in a similar way to (Delgrande and Peppas 2011).

Acknowledgments

We would like to thank the reviewers for their helpful comments.

This work benefited from the support of the project AMANDE ANR-13-BS02-0004 of the French National Research Agency (ANR).

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.
- Amgoud, L., and Cayrol, C. 2002. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence* 34(1-3):197–215.
- Baral, C.; Kraus, S.; Minker, J.; and Subrahmanian, V. S. 1992. Combining knowledge bases consisting of first-order theories. *Computational Intelligence* 8(1):45–71.
- Baral, C.; Kraus, S.; and Minker, J. 1991. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering* 3(2):208–220.
- Baumann, R. 2012. What does it take to enforce an argument? Minimal change in abstract argumentation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'12)*, 127–132.
- Bench-Capon, T. J. M. 2002. Value-based argumentation frameworks. In *Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR'02)*, 443–454.
- Besnard, P., and Doutre, S. 2004. Checking the acceptability of a set of arguments. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR'04)*, 59–64.
- Bisquert, P.; Cayrol, C.; Dupin de Saint-Cyr, F.; and Lagasquie-Schiex, M.-C. 2011. Change in argumentation systems: Exploring the interest of removing an argument. In *Proceedings of the International Conference on Scalable Uncertainty Management (SUM'11)*, 275–288.
- Boella, G.; Kaci, S.; and van der Torre, L. 2009a. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09)*, 107–118.
- Boella, G.; Kaci, S.; and van der Torre, L. 2009b. Dynamics in argumentation with single extensions: attack refinement and the grounded extension. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, 1213–1214.
- Cadoli, M.; Donini, F. M.; Liberatore, P.; and Schaerf, M. 1999. The size of a revised knowledge base. *Artificial Intelligence* 115(1):25–64.
- Caminada, M. 2006. On the issue of reinstatement in argumentation. In *Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA'06)*, 111–123.
- Cayrol, C.; Dupin de Saint-Cyr, 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. *Artificial Intelligence Journal* 171:730–753.
- Delgrande, J. P., and Peppas, P. 2011. Revising horn theories. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, 839–844.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence* 77(2):321–357.
- Dunne, P. E.; Dvořák, W.; Linsbichler, T.; and Woltran, S. 2013. Characteristics of multiple viewpoints in abstract argumentation. In *Proceedings of the 4th Workshop on Dynamics of Knowledge and Belief (DKB-2013)*, 16–30.
- Fagin, R.; Kuper, G. M.; Ullman, J. D.; and Vardi, M. Y. 1986. Updating logical databases. *Advances in Computing Research* 3:1–18.
- Gabriellini, S., and Torroni, P. 2013. MS Dialogues: Persuading and getting persuaded, A model of social network debates that reconciles arguments and trust. In *Proceedings of the Tenth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2013)*.
- Gärdenfors, P. 1988. *Knowledge In Flux*. Cambridge, UK: Cambridge University Press.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.
- Konieczny, S. 2000. On the difference between merging knowledge bases and combining them. In *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR'00)*, 135–144.
- 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 *Proceedings of the 14th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA XIV)*.
- Nebel, B. 1998. How hard is it to revise a belief base? In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, volume 3: Belief Change. Kluwer. 77–145.
- Qi, G.; Liu, W.; and Bell, D. A. 2006. Knowledge base revision in description logics. In *Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA'06)*, 386–398.
- Yager, R. R. 1988. On ordered weighted averaging aggregation operators in multi-criteria decision making. *IEEE Transactions on Systems, Man and Cybernetics* 18:183–190.