How Different Are Two Argumentation Semantics?

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

[1] IRIT, Université Toulouse 1 Capitole, France
[2] Institute of Information Systems, TU Wien, Austria
doutre@irit.fr jmailly@dbai.tuwien.ac.at

Résumé

Ces dernières années, de nombreuses sémantiques pour l'argumentation ont été définies, et leurs propriétés étudiées. Des critères de comparaison entre sémantiques ont été proposés, mais aucune mesure de différence entre sémantique n'a été définie. De telles mesures se révèlent utiles dans le cas où la sémantique associée à un système d'argumentation doit être changée, d'une manière qui assure que la nouvelle sémantique n'est pas trop éloignée de l'ancienne. Trois principales mesures de différence sont proposées dans cet article. Nous montrons que certaines de ces mesures sont des distances, des semi-distances, ou des pseudo-distances.

Abstract

In the last decades, many argumentation semantics have been defined, and their properties studied. Some comparison criteria between semantics have been highlighted, but no measures of difference between semantics have been defined. Such measures turn helpful in the case where the semantics associated to an argumentation framework may have to be changed, in a way that ensures that the new semantics is not too dissimilar from the old one. Three main notions of difference measures between semantics are defined in this paper. Some of these measures are shown to be distances, semi-distances or pseudo-distances.

1 Introduction

Abstract argumentation frameworks (AFs) are classically associated with a semantics which allows to evaluate arguments' statuses, determining sets of jointly acceptable arguments called extensions [10, 2]. In [5, 4], a method to modify an AF in order to satisfy a constraint (a given set of arguments should be an extension, or at least included in an extension) is defined; this process is called extension enforcement. The authors distinguish between conservative enforcement when the semantics does not change (only the AF changes) and liberal enforcement when the semantics changes. But they do not explain why the semantics should change, nor which semantics should be the new one.

Apart from this use of a semantic change for an extension enforcement purpose, a change of the semantics may be necessary for other reasons, for instance, for computational purposes: if a given semantics was appropriate at some point in a certain context for some AF, one may imagine that changes over time on the structure of the AF (number of arguments, of attacks) may make this semantics too "costly" to compute. It may then be interesting to pick up another semantics to apply to the AF, possibly not too dissimilar to the former one.

In another revision context, [9] defines revision operators for AFs which proceeds in two steps. First, revised extensions are computed, then a set of AFs is associated with these revised extensions. Indeed, it is not possible in general to associate a single AF with an arbitrary set of extensions, under a chosen semantics. Modifying the semantics in the revision process could permit to obtain a single AF in some situations, or at least to minimize the number of AFs in the result.

Whatever be the context where a semantic change is necessary, we think that such a semantic change should not be performed any old how, and should respect some kind of minimality, exactly as belief change operations usually require minimal change (see e.g. [15] for belief revision in a propositional setting). Defining *difference measures* between semantics, to quantify how much a semantics is dissimilar to another one, allows to define different minimality criteria. Such criteria can be used to select the new semantics among several options when a semantic change occurs.

Main contribution We propose in this paper three sensible ways to quantify the difference between two semantics:

- depending on the properties which characterize the semantics;
- depending on the relations between semantics;
- depending on the acceptance statuses of arguments the semantics lead to.

The first ones (property-based and relation-based) are said to be *absolute* measures, since they only depend on the considered semantics; they apply to any graph. The last one (acceptance-based) is said to be *relative*: the definition of the measure depends on a particular AF. We study the properties of our measures, in particular we show that some of them are distances, semi-distances or pseudo-distances.

2 Background Notions

An argumentation framework (AF) [10] is a directed graph $\langle A, R\rangle$ where the nodes in A represent abstract entities called *arguments* and the edges in R represent *attacks* between arguments. $(a_i, a_j) \in R$ means that a_i attacks a_j; a_i is called an *attacker* of a_j. We say that an argument a_i (resp. a set of arguments S) defends the argument a_j against its attacker a_k if a_i (resp. any 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 semantics allow to determine which sets of arguments can be collectively accepted [10, 16, 3, 7, 11, 8, 12].

Definition 1. *Let $F = \langle A, R\rangle$ be an AF. A set of arguments $S \subseteq A$ 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 of its arguments against all of their attackers;*
- *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. $\subseteq$);*
- *a stage extension of F if S is conflict-free and there is no conflict-free T such that $S^+_R \subset T^+_R$;*
- *a semi-stable extension of F if S is admissible and there is no admissible T such that $S^+_R \subset T^+_R$;*
- *an ideal set of F if S is admissible and S is included in each preferred extension;*
- *an ideal extension of F if S is a maximal (w.r.t. $\subseteq$) ideal set of F;*
- *an eager extension of F if S is a maximal (w.r.t. $\subseteq$) admissible set that is a subset of each semi-stable extension.*

These semantics are denoted, respectively, $cf, adm, na, co, pr, st, gr, stg, sem, is, id, eg$. For each σ of them, $Ext_\sigma(F)$ denotes the set of σ-extensions of F.

Let us recall the definition of some usual decision problems for argumentation.

Definition 2. *Let $F = \langle A, R\rangle$ be an AF and σ a semantics.*

- *An argument $a_i \in A$ is said to be* credulously accepted *by F w.r.t. σ if $\exists E \in Ext_\sigma(F)$ s.t. $a_i \in E$.*
- *An argument $a_i \in A$ is said to be* skeptically accepted *by F w.r.t. σ if $\forall E \in Ext_\sigma(F), a_i \in E$.*

The set of credulously (resp. skeptically) accepted arguments in F w.r.t. σ is denoted $cr_\sigma(F)$ (resp. $sk_\sigma(F)$).

Example 1. *Let us consider the argumentation framework F_1 given at Figure 1, and let us illustrate some of the semantics, and related decision problems. $Ext_{adm}(F_1) = \{\emptyset, \{a_1\}, \{a_4\}, \{a_4, a_6\}, \{a_1, a_4, a_6\}, \{a_1, a_3\}, \{a_1, a_4\}\}$, $Ext_{st}(F_1) = \{\{a_1, a_4, a_6\}\}$, $Ext_{pr}(F_1) = \{\{a_1, a_4, a_6\}, \{a_1, a_3\}\}$, $Ext_{co}(F_1) = \{\{a_1, a_4, a_6\}, \{a_1, a_3\}, \{a_1\}\}$, $Ext_{gr}(F_1) = \{\{a_1\}\}$. a_1 is skeptically accepted in F_1 w.r.t. the stable, preferred, complete and grounded semantics. a_4 is credulously accepted in F_1 w.r.t. the preferred and complete semantics, but it is not w.r.t. the grounded semantics.*

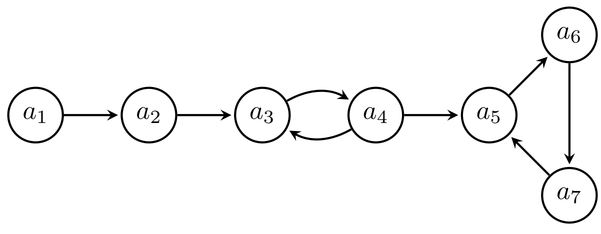

FIGURE 1 – The AF F_1

In order to compare, in the following section, the semantics, and propose measures of their differences, let us introduce a useful notation: given two sets X, Y, $X\Delta Y$ is the symmetric difference between X and Y. Let us recall also the definition of a distance and an aggregation function.

Definition 3. *Given a set E, a mapping d from $E \times E$ to $\mathbb{R}^+$ is*
- *a* pseudo-distance *if it satisfies weak coincidence, symmetry and triangular inequality;*
- *a* semi-distance *if it satisfies coincidence and symmetry;*
- *a* distance *if it satisfies coincidence, symmetry and triangular inequality.*

weak coincidence $\forall x \in E, d(x,x) = 0$;

coincidence $\forall x, y \in E, d(x,y) = 0$ *iff* $x = y$;

symmetry $\forall x, y \in E, d(x,y) = d(y,x)$;

triangular inequality $\forall x, y, z \in E, d(x,y) + d(y,z) \geq d(x,z)$.

Definition 4. *An* aggregation function *is a function $\otimes$ which associates a non-negative real number to every finite tuple of non-negative numbers, and which satisfies:*

non-decreasingness *if* $y \leq z$ *then* $\otimes(x_1, \dots, y, \dots, x_n) \leq \otimes(x_1, \dots, z, \dots, x_n)$;

minimality $\otimes(x_1, \dots, x_n) = 0$ *iff* $x_1 = \dots = x_n = 0$;

identity $\forall x \in \mathbb{R}^+, \otimes(x) = x$.

For instance, we will use the sum $\sum$ as an aggregation function.

3 Property-based Difference Measures

We propose a first way to measure how much two semantics are different. Here, the idea is to split a semantics into a set of properties which characterize it, and to give a weight to each property, these weights corresponding to the importance of the property in the context where the semantics have to be compared. Then, measuring the difference between two semantics is equivalent to adding the weight of the properties which appear in the characterization of exactly one of the semantics.

Definition 5. *A set of properties $\mathcal{P}$ characterizes a semantics σ if for each AF F,*

1. *each σ-extension of F satisfies each property from $\mathcal{P}$,*
2. *each set of arguments which satisfies each property from $\mathcal{P}$ is a σ-extension of F,*
3. *$\mathcal{P}$ is a minimal set (w.r.t $\subseteq$) among those which satisfy* 1. *and* 2.

$Prop(\sigma)$ denotes the set of properties that characterizes a semantics σ.

Beyond the use of characterizations to define difference measures, characterizations can have a computational interest. For instance, verifying if a set of arguments is a σ-extension can be done by checking if it satisfies all the properties in $Prop(\sigma)$. In this case, the computation can stop as soon as one of the properties is not satisfied.

Let us point out interesting properties, and establish which ones characterize each semantics. We distinguish between absolute properties (which concern only a set of arguments itself, Definition 6) and relative properties (which concern a set of arguments with respect to other sets of arguments, Definition 7).

Definition 6. *Given an AF $F = \langle A, R \rangle$, a set of arguments S satisfies*
- *conflict-freeness if S is conflict-free;*
- *acceptability if S defends itself against each attacker;*
- *reinstatement if S contains all the arguments that it defends;*
- *complement attack if each argument in $A \backslash S$ is attacked by S.*

Definition 7. *Given an AF $F = \langle A, R \rangle$ and a set of properties $\mathcal{P}$, a set of arguments S satisfies*
- *$\mathcal{P}$-maximality if S is maximal (w.r.t. $\subseteq$) among the sets of arguments which satisfy $\mathcal{P}$;*
- *$\mathcal{P}$-minimality if S is minimal (w.r.t. $\subseteq$) among the sets of arguments which satisfy $\mathcal{P}$;*
- *$\mathcal{P}$-inclusion if S is included in each set of arguments which satisfies $\mathcal{P}$;*
- *$\mathcal{P}$-R-maximality if S has a maximal range (w.r.t. $\subseteq$) among the sets of arguments which satisfy $\mathcal{P}$.*

It can be noticed that, by definition, if a set S satisfies $\mathcal{P}$-maximality (resp. $\mathcal{P}$-minimality, $\mathcal{P}$-R-maximality), then S satisfies $\mathcal{P}$.

Now, we establish a characterization of the different semantics, that follows from the previous definitions.

Proposition 1. *The extension-based semantics considered in this paper can be characterized as follows:*
- $Prop(cf) = \{$*conflict-freeness*$\}$.
- $Prop(adm) = Prop(cf) \cup \{$*acceptability*$\}$.
- $Prop(na) = Prop(cf)$*-maximality.*
- $Prop(co) = Prop(adm) \cup \{$*reinstatement*$\}$.
- $Prop(gr) = Prop(co)$*-minimality.*
- $Prop(pr) = Prop(adm)$*-maximality.*
- $Prop(sem) = Prop(adm)$*-R-maximality.*
- $Prop(stg) = Prop(cf)$*-R-maximality.*
- $Prop(st) = Prop(cf) \cup \{$*complement attack*$\}$.
- $Prop(is) = Prop(adm) \cup \{Prop(pr)$*-inclusion*$\}$.
- $Prop(id) = Prop(is)$*-maximality.*
- $Prop(eg) = Prop(pr) \cup \{Prop(sem)$*-inclusion*$\}$.

Let us notice that we could consider other properties, and give alternative characterizations of the semantics. Even if the value of the difference between two semantics (obviously) depends of the chosen characterizations, the general definition of property-based difference measures is the same whatever the characterizations.

Our intuition which leads to define the characterization as the minimal set of properties is related to computational issues. Indeed, computing some reasoning tasks related to the semantics thanks to the semantics characterization can be done more efficiently with this definition. For instance, to determine whether a set of arguments is a stable extension of a given AF, checking the satisfaction of conflict-freeness and complement attack proves enough. For instance, we could add $Prop(adm)$-maximality in the characterization of the stable semantics, but computing the result of our problem would then be harder.

A weight can be associated to each property, depending on the importance of the property in a certain context.

Definition 8. *Let $\mathcal{P}$ be a set of properties. Let w be a function which maps each property $p \in \mathcal{P}$ to a strictly positive real number $w(p)$. Given σ_1, σ_2 two semantics such that $Prop(\sigma_1) \subseteq \mathcal{P}$ and $Prop(\sigma_2) \subseteq \mathcal{P}$, the* property-based difference measure *δ^w_{prop} between σ_1 and σ_2 is defined as:*

$$\delta^w_{prop}(\sigma_1, \sigma_2) = \sum_{p_i \in Prop(\sigma_1) \Delta Prop(\sigma_2)} w(p_i)$$

The specific property-based difference measure defined when all the properties have the same importance is defined as follows.

Definition 9. *Given two semantics σ_1, σ_2, the property-based difference measure δ_{prop} is defined by $\delta_{prop}(\sigma_1, \sigma_2) = |Prop(\sigma_1) \Delta Prop(\sigma_2)|$.*

Example 2. *Let us suppose that the initial semantics is the admissible one. When we consider δ_{prop}, naive and preferred semantics are "equivalent", since $\delta_{prop}(adm, na) = \delta_{prop}(adm, pr) = 3$. On the opposite, with a weighted measure δ^w_{prop} such that $w(Prop(cf)$-maximality$) = 1$ and $w(Prop(adm)$-maximality$) = 2$, the naive semantics is "better" since $\delta_{prop}(adm, na) < \delta_{prop}(adm, pr)$.*

Proposition 2. *Given a set of semantics $\mathcal{S}$, the property-based measures defined on $\mathcal{S}$ are distances.*

4 Relation-based Difference Measures

The second absolute method to measure the difference between semantics that we propose, is based on the fact that most of the usual semantics are related according to some notions. For instance, it is well-known that each preferred extension of an AF is also a complete extension of it, and the grounded extension is also complete, but in general it is not a preferred extension. The preferred semantics may thus be seen closer to the complete semantics, than to the grounded semantics. We formalize this idea with the notion of semantics relation graph.

Definition 10. *Let $\mathcal{S} = \{\sigma_1, \ldots, \sigma_n\}$ a set of semantics. A* semantics relation graph *on $\mathcal{S}$ is defined by $Rel(\mathcal{S}) = \langle \mathcal{S}, D \rangle$ with $D \subseteq \mathcal{S} \times \mathcal{S}$.*

This abstract notion of relation graph, where the nodes are semantics, can be instantiated with the inclusion relation between the extensions of an AF.

Definition 11. *Let $\mathcal{S} = \{\sigma_1, \ldots, \sigma_n\}$ a set of semantics. The* extension inclusion graph *of $\mathcal{S}$ is defined by $Inc(\mathcal{S}) = \langle \mathcal{S}, D \rangle$ with $D \subseteq \mathcal{S} \times \mathcal{S}$ 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 $\sigma_k \in \mathcal{S}$ $(k \neq i, k \neq j)$ such that for each AF F, $Ext_{\sigma_i}(F) \subseteq Ext_{\sigma_k}(F)$ and $Ext_{\sigma_k}(F) \subseteq Ext_{\sigma_j}(F)$.*

This idea is discussed in [2], but that paper does not formalize the notion of relation between semantics as we do here.

Example 3. *For instance, when $\mathcal{S} = \{co, pr, st, gr, stg, sem, is, id, eg, adm, cf, na\}$, $Inc(\mathcal{S})$ is the graph given at Figure 2.*

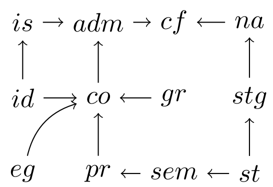

FIGURE 2 – Extension Inclusion Graph $Inc(\mathcal{S})$

Now, we define a family of difference measures between semantics which is based on the semantics relation graphs, to measure what it costs for an agent to change her semantics.

Definition 12. *Given $\mathcal{S}$ a set of semantics, a $\mathcal{S}$-*relation difference measure *is the mapping from two semantics $\sigma_1, \sigma_2 \in \mathcal{S}$ to the non-negative integer $\delta_{Rel,\mathcal{S}}(\sigma_1, \sigma_2)$ which is the length of the shortest non-oriented path between σ_1 and σ_2 in $Rel(\mathcal{S})$.*
In particular, the $\mathcal{S}$-inclusion measure is the length of the shortest non-oriented path between σ_1 and σ_2 in $Inc(\mathcal{S})$, denoted by $\delta_{Inc,\mathcal{S}}(\sigma_1, \sigma_2)$.

Example 4. *Given two semantics σ_1 and σ_2 which are neighbours in the graph given at Figure 2, the difference measure $\delta_{Inc,\mathcal{S}}(\sigma_1, \sigma_2)$ is obviously 1. Otherwise, if several paths allow to reach σ_2 from σ_1, then* *the difference is the length of the minimal one. For instance, $\delta_{dep,\mathcal{S}}(st, cf) = 3$ since the minimal path is $st \to stg \to na \to cf$, but other paths exist (for instance, $st \to sem \to pr \to co \to adm \to cf$). Since here the question is to define the difference between semantics, the possibility to obtain several minimal paths (for instance, there are two minimal paths between the ideal and admissible semantics: $id \to is \to adm$ and $id \to co \to adm$) is not problematic.*

Proposition 3. *The $\mathcal{S}$-inclusion difference measure is a distance.*

We could instantiate the relation graph with another relation between semantics. For instance, we can define a graph such that there is an edge from σ_1 to σ_2 if $\forall F, sk_{\sigma_1}(F) \subseteq sk_{\sigma_2}(F)$, or similarly with the credulous acceptance. This requires a deep investigation of the relations between semantics w.r.t. skeptical or credulous acceptance.

For the possible instantiations of the relation graph that have been proposed, we can also define a relative version. In this case, the edges in the graph depend on the inclusion relations for a given AF, while our first proposal considers the inclusion relations which are true for any AF. This AF-based relation graph can lead to an interesting new measure.

The graph resulting from the intertranslatability relationship of semantics [13] may also provide another instantiation of the relation graph.

5 Acceptance-based Difference Measures

We have defined previously two approaches to quantify the difference between semantics which are absolute, which means that the difference between two semantics is always the same, whatever the situation and the AF. It may be interesting for some applications to take into account the AF of the agent to measure the difference between the semantics. We propose here such a family of measures. Now, the difference between semantics depends on the acceptance status of arguments in a given AF, w.r.t. the different semantics in consideration.

Our first acceptance-based measure quantifies the difference between the σ_1-extensions and the σ_2-extension of the AF to quantify the difference between σ_1 and σ_2.

Definition 13. *Let F be an AF, d be a distance between sets of arguments, and $\otimes$ be an aggregation function. The* F-d-$\otimes$-extension-based difference measure *$\delta_F^{d,\otimes}$ is defined by $\delta_F^{d,\otimes}(\sigma_1, \sigma_2) = \otimes_{\epsilon \in Ext_{\sigma_1}(F)} \min_{\epsilon' \in Ext_{\sigma_2}(F)} d(\epsilon, \epsilon')$.*

Proposition 4. *In general, the extension-based difference measures are not distances, they do not satisfy coincidence, symmetry.*

Example 5. *For instance, we consider the Hamming distance between sets of arguments, defined as $d_H(s_1, s_2) = |s_1 \Delta s_2|$. Now, we define the F_1-d_H-$\sum$-extension-based difference measure $\delta_F^{d_H,\sum}$ from d_H and the AF F_1 given at Figure 1. Its set of stable extensions is $Ext_{st}(F_1) = \{\{a_1, a_4, a_6\}\}$.*

When measuring the difference between the stable semantics and the other classical Dung's semantics, we obtain:

- *$\delta_{F_1}^{d_H,\sum}(st, gr) = 2$ since $Ext_{gr}(F_1) = \{\{a_1\}\}$;*
- *$\delta_{F_1}^{d_H,\sum}(st, pr) = 0$ since $Ext_{pr}(F_1) = \{\{a_1, a_3\}, \{a_1, a_4, a_6\}\}$; on the opposite, $\delta_{F_1}^{d_H}(pr, st) = 3$;*
- *$\delta_{F_1}^{d_H,\sum}(st, co) = 0$ since $Ext_{co}(F_1) = \{\{a_1\}, \{a_1, a_3\}, \{a_1, a_4, a_6\}\}$.*

The following result shows that the restriction of the extension-based measure to some particular sets of semantics leads to satisfy the coincidence property.

Proposition 5. *For a given F and a given set of semantics $\mathcal{S} = \{\sigma_1, \dots, \sigma_n\}$, if for all $\sigma_i, \sigma_j \in \mathcal{S}$ such that $\sigma_i \neq \sigma_j$, $Ext_{\sigma_i}(F) \not\subseteq Ext_{\sigma_j}(F)$, then the extension-based measure $\delta_F^{d_H,\sum}$ satisfies coincidence.*

Even in this case, the measure does no satisfy all the properties of distances. However, we can use the intuition behind this measure to define another one.

Definition 14. *Let F be an AF, d be a distance between sets of arguments, and $\otimes$ be an aggregation function. The* symmetric F-d-$\otimes$-extension-based difference measure *$\delta_{F,sym}^{d,\otimes}$ is defined by $\delta_{F,sym}^{d,\otimes}(\sigma_1, \sigma_2) = \max(\delta_F^{d,\otimes}(\sigma_1, \sigma_2), \delta_F^{d,\otimes}(\sigma_2, \sigma_1))$.*

This measure satisfies the distance properties under some conditions.

Proposition 6. *For a given F and a given set of semantics $\mathcal{S} = \{\sigma_1, \dots, \sigma_n\}$, if for all $\sigma_i, \sigma_j \in \mathcal{S}$ such that $\sigma_i \neq \sigma_j$, $Ext_{\sigma_i}(F) \neq Ext_{\sigma_j}(F)$, then the symmetric extension-based measure $\delta_{F,sym}^{d_H,\sum}$ is a semi-distance.*

We propose here some extension-based measures. We can also use the set of skeptically (resp. credulously) accepted arguments instead of the whole set of extensions to define a difference measure between semantics.

Definition 15. *Given F an AF, d a distance between sets of arguments, and $\mathcal{S}$ a set of semantics, the F-d-*skeptical acceptance difference measure *$\delta_{F,sk}^d$ is defined, for any $\sigma_1, \sigma_2 \in \mathcal{S}$, by*

$$\delta_{F,sk}^d(\sigma_1, \sigma_2) = d(sk_{\sigma_1}(F), sk_{\sigma_2}(F))$$

The F-d-credulous acceptance difference measure $\delta^d_{F,sk}$ *is defined, for any* $\sigma_1, \sigma_2 \in \mathcal{S}$*, by*

$$\delta^d_{F,cr}(\sigma_1, \sigma_2) = d(cr_{\sigma_1}(F), cr_{\sigma_2}(F))$$

If two semantics lead to the same set of credulously (resp. skeptically) accepted arguments, then these measures cannot distinguish between these semantics. Other properties are satisfied.

Proposition 7. *Given* F *and* AF *and* d *a distance, the* F-d*-skeptical acceptance difference measure and the* F-d*-credulous acceptance difference measure are pseudo-distances.*

6 Combining Measures to Obtain New Minimality Criteria

In the context of a semantic change, our difference measures may be used to define different minimality criteria. With σ_1 the initial semantics, and $\mathcal{S}$ the set of options for the new semantics (for instance, the semantics which permit to have a single AF as the result of a revision, as mentioned in the introduction), the new semantics should be $\sigma' \in \mathcal{S}$ such that $\forall \sigma'' \in \mathcal{S}$, $\delta(\sigma_1, \sigma') \leq \delta(\sigma_1, \sigma'')$, with δ the chosen measure. However, this does not always lead to a single result, as we have seen in the previous examples. So, to distinguish between several possible semantics which are minimal with respect to a first measure, we can apply a second measure to refine the result; and so on. In general, there is no guarantee to obtain a single result, but using different levels of minimality, induced by different measures, permits the agent to choose her new semantics amongst fewer possible options. Also, the order of application of the different measures may lead to different results.

Definition 16. *Let* $D = \langle \delta_1, \ldots, \delta_n \rangle$ *a vector of difference measures between semantics. Let* σ *be a semantics, and* $\mathcal{S}$ *a set of semantics. The* D-minimal semantic change selection function *is defined by* $\gamma_D(\sigma, \mathcal{S}) = \gamma^n_D(\sigma, \mathcal{S})$ *with* γ^n_D *as follows:*

$$\begin{aligned} \gamma^1_D(\sigma, \mathcal{S}) &= \{\sigma_i \in \mathcal{S} \mid \forall \sigma_j \in \mathcal{S}, \delta_1(\sigma, \sigma_i) \leq \delta_1(\sigma, \sigma_j)\} \\ \gamma^k_D(\sigma, \mathcal{S}) &= \{\sigma_i \in \gamma^{k-1}_D(\sigma, \mathcal{S}) \mid \forall \sigma_j \in \gamma^{k-1}_D(\sigma, \mathcal{S}), \\ &\quad \delta_k(\sigma, \sigma_i) \leq \delta_k(\sigma, \sigma_j)\} \end{aligned}$$

Let us notice that this definition is general enough to encompass any difference measure yet to be defined.

7 Conclusion

In this paper, we have defined several ways to quantify the difference between extension-based semantics. Some of them are absolute (they only depend on the semantics), while the other ones are relative (they depend on the considered AF). Let us mention the fact that there is no general relation between these difference measures; for instance we have seen on several examples that it may occur that $\delta_1(\sigma_1, \sigma_2) > \delta_1(\sigma_1, \sigma_3)$ while $\delta_2(\sigma_1, \sigma_2) < \delta_2(\sigma_1, \sigma_3)$. When a semantic change occurs, this permits the agent to use some very different notions of minimality to select the new semantics, depending on which difference measures make sense in the context of her application. In addition, the combination of these "basic" measures permits to express even more notions of minimality.

Let us notice that only the relation-based and property-based measures are distances, other methods failing in general to satisfy the distance properties, which seem to be desirable to quantify the difference between objects. However, the skeptical and credulous acceptance difference measures are pseudo-distances. Further study could lead to identify the necessary conditions that a set of semantics must satisfy to ensure that these are distances.

	δ^w_{prop}	$\delta_{Inc,\mathcal{S}}$	$\delta^{d,\sum}_F$	$\delta^{d,\sum}_{F,sym}$	$\delta^d_{F,sk}$	$\delta^d_{F,cr}$
WC	✓	✓	∘	✓	✓	✓
Co	✓	✓	×	∘	×	×
Sym	✓	✓	×	✓	✓	✓
TI	✓	✓			✓	✓

Table 1 – Summary of properties satisfied by our measures

Table 1 depicts the properties satisfied by our measures. WC, Co, Sym and TI stand respectively for weak coincidence, coincidence, symmetry and triangular inequality. A ✓ symbol means that the property is always satisfied, and × means that it is not satisfied in general. ∘ means that the property is satisfied under some additional assumption.

We consider several tracks for future works. We have noticed that we can order semantics, with respect to an initial semantics σ and a measure δ: $\sigma_1 \leq_{\sigma,\delta} \sigma_2$ if and only if $\delta(\sigma, \sigma_1) \leq \delta(\sigma, \sigma_2)$. In this case, we can investigate the relation of the orderings defined by different measures. For instance, if some pairs (σ, δ_1) and (σ, δ_2) lead to the same ordering, then we can choose to use the measure which is the least expensive one to compute among δ_1 and δ_2.

We also plan to define a similar notion of difference measures for labelling-based semantics [2], and for ranking-based semantics [1, 14, 6]. In this last context, we need to determine whether some relevant properties characterize the ranking which is used to evaluate arguments, or to determine meaningful notions of difference between the rankings.

Finally, we will investigate the question which is mentioned in the introduction: using (minimal) semantic change to define enforcement and revision methods. In particular, we think that semantic change can be used to guarantee minimal change on the attack relation, or to ensure that the result of the process is a single AF.

Acknowledgements

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

A Proofs

Proof of Proposition 2. From our definition of characterizations, the mapping that associates a semantics σ to a set of properties $Prop(\sigma)$ guarantees that a semantics cannot be associated with two different sets of properties, and a same set of properties cannot correspond to different semantics.

The weighted sum on sets of properties obviously defines a distance (in particular, when all weights are identical, we obtain the well-known Hamming distance; other weights just define generalization of Hamming distance). Since we can identify the semantics to the sets of properties, δ^w_{prop} is a distance. □

Proof of Proposition 3. From the definition of the Σ-relation graph,

- the difference between σ_1 and σ_2 is 0 iff they are the same node of the graph (i.e. $\sigma_1 = \sigma_2$), so coincidence is satisfied;
- the shortest path between two semantics σ_1, σ_2 has the same length whatever the direction of the path (from σ_1 to σ_2, or vice-versa), since we do not consider the direction of arrows, so symmetry is satisfied;
- the shortest path between σ_1 and σ_3 is at worst the concatenation of the paths $(\sigma_1, \ldots, \sigma_2)$ and $(\sigma_2, \ldots, \sigma_3)$, or (if possible) a shorter one, so triangular inequality is satisfied.

□

Proof of Proposition 4. Example 5 gives the counter-examples for coincidence and symmetry. □

Proof of Proposition 5. We consider a given AF F and a set of semantics $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$, such that for all $\sigma_i, \sigma_j \in \Sigma$ with $\sigma_i \neq \sigma_j$, $Ext_{\sigma_i}(F) \nsubseteq Ext_{\sigma_n}(F)$.

Obviously, for any semantics σ_i, $\delta^{d_H,\Sigma}_F(\sigma_i, \sigma_i) = 0$. Now, let us assume the existence of two semantics $\sigma_i, \sigma_j \in \Sigma$ such that $\delta^{d_H,\Sigma}_F(\sigma_i, \sigma_j) = 0$. We just rewrite this, following the definition of the measure: $\sum_{\epsilon \in Ext_{\sigma_i}(F)} \min_{\epsilon' \in Ext_{\sigma_j}(F)} d_H(\epsilon, \epsilon') = 0$. Since all distances are non-negative number, if the sum is equal to zero it means that $\forall \epsilon \in Ext_{\sigma_i}(F)$, $\min_{\epsilon' \in Ext_{\sigma_j}(F)} d_H(\epsilon, \epsilon') = 0$. Because of the properties of the Hamming distance, it means that $\epsilon \in Ext_{\sigma_j}$, and so $Ext_{\sigma_i} \subseteq Ext_{\sigma_j}$. From our starting assumption, we deduce that $\sigma_i = \sigma_j$. □

Proof of Proposition 6. From the definition of the measure, $\delta^{d_H,\Sigma}_{F,sym}(\sigma_1, \sigma_2) = 0$ iff $Ext_{\sigma_1}(F) = Ext_{\sigma_2}(F)$. Under our assumptions, this is possible only if $\sigma_1 = \sigma_2$. The other direction is trivial, so coincidence is satisfied. Symmetry is obviously satisfied, since σ_1, σ_2 can be inverted in $\max(\delta^{d,\otimes}_F(\sigma_1, \sigma_2), \delta^{d,\otimes}_F(\sigma_2, \sigma_1))$. □

Proof of Proposition 7. Weak coincidence and symmetry are trivial from the definition of the measures.

$\delta^d_{F,sk}(\sigma_1, \sigma_2) + \delta^d_{F,sk}(\sigma_2, \sigma_3)$
$= d(sk_{\sigma_1}(F), sk_{\sigma_2}(F)) + d(sk_{\sigma_2}(F), sk_{\sigma_3}(F))$
$\geq d(sk_{\sigma_1}(F), sk_{\sigma_3}(F)) = \delta^d_{F,sk}(\sigma_1, \sigma_3)$

The same reasoning apply for the credulous acceptance measure. So both satisfy the triangular inequality. Coincidence is not satisfied by the skeptical acceptance measure. For instance, for each AF F, $\emptyset \in Ext_{cf}(F)$ and $\emptyset \in Ext_{adm}(F)$, so $sk_{cf}(F) = sk_{adm}(F) = \emptyset$, and so $\delta^d_{F,skep}(cf, adm) = 0$. The same conclusion holds as soon as two semantics yield the same skeptically or credulously accepted arguments. □

Références

[1] Leila Amgoud and Jonathan Ben-Naim. Ranking-based semantics for argumentation frameworks. In *Proc. SUM'13*, pages 134–147, 2013.

[2] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Eng. Review*, 2011.

[3] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artif. Intell.*, 2005.

[4] Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *Proc. ECAI'12*, pages 127–132, 2012.

[5] Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proc. COMMA'10*, pages 75–86, 2010.

[6] Elise Bonzon, Jérôme Delobelle, Sébastien Konieczny, and Nicolas Maudet. A comparative study of ranking-based semantics for abstract argumentation. In *Proc. AAAI'16*, 2016.

[7] Martin Caminada. Semi-stable semantics. In *Proc. COMMA'06*, 2006.

[8] Martin Caminada. Comparing two unique extension semantics for formal argumentation: Ideal and eager. 2007.

[9] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In *Proc. KR'14*, 2014.

[10] Phan Minh Dung. 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.

[11] Phan Minh Dung, Paolo Mancarella, and Francesca Toni. Adialectic procedure for sceptical, assumption-based argumentation. In *COMMA'06*, 2006.

[12] Paul E. Dunne. The computational complexity of ideal semantics. *Artif. Intell.*, pages 1559–1591, 2009.

[13] Wolfgang Dvorák and Christof Spanring. Comparing the expressiveness of argumentation semantics. In *Proc. COMMA'12*, volume 245, pages 261–272. IOS Press, 2012.

[14] Davide Grossi and Sanjay Modgil. On the graded acceptability of arguments. In *Proc. IJCAI'15*, pages 868–874, 2015.

[15] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[16] Bart Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In *Proc. BNAIC'96*, 1996.