Axiomatic Foundations of Acceptability Semantics

Leila Amgoud **Jonathan Ben-Naim**
IRIT – CNRS
118 route de Narbonne
F-31062 Toulouse Cedex 9, France

Abstract

An argument is a reason or justification of a claim. It has an intrinsic strength and may be attacked by other arguments. Hence, the evaluation of its *overall strength* becomes mandatory, especially for judging the reliability of its claim. Such an evaluation is done by acceptability semantics.

The aim of this paper is to set up the *foundations* of acceptability semantics. Foundations are important not only for a better understanding of the evaluation process in general, but also for clarifying the basic assumptions underlying semantics, for comparing different (families of) semantics and identifying families of semantics that have not been explored yet.

The paper defines the building blocks of a semantics. It introduces key *concepts* and *principles* on which an evaluation is based. Each concept (principle) is described by an axiom. We investigate properties of semantics that satisfy the axioms, show the foundations of the two crucial notions of reinstatement and defence, and analyse some existing semantics against the axioms.

Introduction

An *argument* gives reason to support a claim that is questionable, or open to doubt. It is made of three components: *premises* representing the reason, a *conclusion* which is the supported claim, and a *link* showing how the premises lead to the conclusion. The link is hence the logical "glue" that binds premises and conclusions together.

An argument has an *intrinsic strength* which may come from different sources: the certainty degree of its reason (Amgoud and Cayrol 2002), the importance of the value it promotes if any (Bench-Capon 2003), the reliability of its source (Parsons et al. 2011), ... Whatever its intrinsic strength (strong or weak), an argument may be *attacked* by other arguments. An attack amounts to undermining one of the components of an argument, and has thus a negative impact on its target. An evaluation of the *overall strength* (or *overall acceptability*) of an argument becomes mandatory, namely for judging whether or not its conclusion is reliable.

The evaluation of arguments has received great interest from the computational argumentation community. Indeed, two families of acceptability semantics were defined for this purpose: *extension* semantics and *gradual* semantics. Extension semantics were initially introduced by Dung (1995). Starting with a set of arguments and attacks between them, they return a set of extensions, each of which is a set of arguments that are acceptable together. Then, using a membership criterion, a qualitative acceptability degree is assigned to each argument. Examples of such semantics are the classical semantics of Dung (complete, stable, preferred, ...) and their different refinements (e.g. (Baroni, Giacomin, and Guida 2005; Caminada 2006b; Grossi and Modgil 2015; Dung, Mancarella, and Toni 2007)). Unlike extension semantics, gradual semantics do not compute extensions. They assign a numerical acceptability degree to each argument. Examples of such semantics are h-Categorizer (Besnard and Hunter 2001), Bbs, Dbs (Amgoud and Ben-Naim 2013) and those proposed in (Matt and Toni 2008; Bonzon, Maudet, and Moretti 2014).

Despite the great interest in semantics, there are only a few works on *foundations* of semantics. Baroni and Giacomin (2007) defined axioms that a semantics would satisfy. However, those axioms are only suited for extension semantics. Furthermore, they are mainly properties of extensions and not of overall strengths of arguments. Finally, most of the axioms are based on concepts (like defence and reinstatement) whose own foundations are unclear. Amgoud and Ben-Naim (2013) proposed another set of axioms for the family of gradual semantics. The axioms are on the ranking of arguments with regard to their overall strengths. While some of the axioms (like independence and abstraction) are primitive, others are much more complex (like counter-transitivity) and their own foundations need to be clarified.

Hence, existing axiomatic studies do not tell much on the foundations of acceptability semantics. Foundations are important not only for a better understanding of the evaluation process, but also for comparing semantics and identifying families of semantics that have not been explored yet.

The aim of this paper is to set up the *foundations* of acceptability semantics. It defines elementary *concepts* and *principles* on which an evaluation of arguments is based. The approach followed in the paper is axiomatic. We introduce a set of axioms, each of which describes a concept or a principle. The axioms are primitive, in that they cannot be decomposed into other axioms. We investigate the properties of semantics that satisfy the axioms. We show in particular the foundations of defence and reinstatement, two key notions of extension semantics. Finally, we analyse existing semantics against the axioms, namely extension semantics proposed by Dung (1995) and the gradual h-Categorizer semantics proposed by Besnard and Hunter (2001). This analysis allows not only a better understanding of the assumptions and choices made by those semantics, but also a clear comparison between semantics of the same family, and between extension semantics and gradual ones.

The paper is structured as follows: We start by introducing some background on argumentation, then we present our list of axioms. Next, we investigate the links between axioms, and show that general versions of some axioms follow from the list of axioms. We then investigate the properties of semantics that satisfy the axioms. Finally, we analyse existing semantics with regard to the axioms.

Basic concepts

An *argumentation framework*, called also *argumentation graph* in the paper, is a set of arguments and a binary relation representing attacks between the arguments. Arguments are abstract entities whose internal structure is not specified.

Definition 1 (Argumentation graph) *An* argumentation graph *is an ordered pair* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, where* $\mathcal{A}$ *is a finite set of arguments and* $\mathcal{R}$ *is a binary relation on* $\mathcal{A}$*, i.e.,* $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$*. Intuitively,* $a\mathcal{R}b$ *means* a attacks b*.*

We present next the list of all notations used in the paper.

Notations: Let $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation graph and $a \in \mathcal{A}$. We denote by $\mathtt{Att}_\mathbf{A}(a)$ the set of all *attackers* of a in $\mathbf{A}$ (i.e. $\mathtt{Att}_\mathbf{A}(a) = \{b \in \mathcal{A} \mid b\mathcal{R}a\}$), by $\mathtt{Att}+_\mathbf{A}(a)$ the set of arguments attacked by a (i.e. $\mathtt{Att}+_\mathbf{A}(a) = \{b \in \mathcal{A} \mid a\mathcal{R}b\}$), and by $\mathtt{Def}_\mathbf{A}(a)$ the set of all *defenders* of a in $\mathbf{A}$ (i.e. $\mathtt{Def}_\mathbf{A}(a) = \{b \in \mathcal{A} \mid \exists\, c \in \mathcal{A}, c\mathcal{R}a$ and $b\mathcal{R}c\}$). For any two argumentation graphs $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{A}' = \langle \mathcal{A}', \mathcal{R}' \rangle$, $\mathbf{A} \oplus \mathbf{A}'$ denotes the argumentation graph $\langle \mathcal{A} \cup \mathcal{A}', \mathcal{R} \cup \mathcal{R}' \rangle$.

Before presenting our definition of semantics, let us first introduce the concept of weighting.

Definition 2 (Weighting) *A* weighting *on a set* X *is a function from* X *to* $[0, 1]$.

We define an acceptability semantics as a function assigning for every argument in an argumentation graph a value between 0 and 1. This value, called *acceptability degree*, represents the *overall strength* of an argument, i.e., the strength issued from the *aggregation* of the intrinsic strength of the argument and the overall strengths of its attackers. The interval $[0, 1]$ may be replaced by any linearly ordered scale.

Definition 3 (Acceptability Semantics) *An* acceptability semantics *is a function* $\mathbf{S}$ *transforming any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *into a weighting on* $\mathcal{A}$*. For* $a \in \mathcal{A}$*,* $\mathtt{Deg}^\mathbf{S}_\mathbf{A}(a)$ *denotes the image of* a *by* $\mathbf{S}(\mathbf{A})$ *and is called* acceptability degree *of* a*.*

Remark: Arguments that get value 1 are very strong while arguments that get value 0 are too weak that one cannot rely on the claims they support, and will be called rejected arguments throughout the paper.

Let us now recall two notions that are useful for the rest of the paper: *isomorphism* and *elementary cycle*.

Definition 4 (Isomorphism) *Let* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *and* $\mathbf{A}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ *be two argumentation graphs. An* isomorphism *from* $\mathbf{A}$ *to* $\mathbf{A}'$ *is a bijective function* f *from* $\mathcal{A}$ *to* $\mathcal{A}'$ *such that the following holds:*

$$\forall\, a, b \in \mathcal{A},\ a\mathcal{R}b \textit{ iff } f(a)\mathcal{R}'f(b).$$

Definition 5 (Elementary Cycle) *An argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, with* $\mathcal{A} = \{a_0, \dots, a_n\}$*, is an* elementary cycle *iff for all* $i \in \{0, 1, 2, \dots, n-1\}$*,* $\mathtt{Att}_\mathbf{A}(a_i) = \{a_{i+1}\}$*, and* $\mathtt{Att}_\mathbf{A}(a_n) = \{a_0\}$*.*

Throughout the paper, we make the following smoothening assumption.

Assumption (Smoothening): All the arguments of an argumentation graph have the same intrinsic strength, assumed to be equal to 1. The choice of value 1 is related to the maximal value of an acceptability degree. The basic idea is that non-attacked arguments keep their intrinsic strength. Unlike in (Dunne et al. 2011), the attacks of any argumentation graph have the same weight.

The reader might wonder why we put emphasis on this smoothening assumption. The reason is that we want to make clearer the assumptions our axioms are based on, and that will be lifted in future work with richer argumentation structures. It is also worthy to say that the smoothening assumption is made in all works in which semantics were defined (e.g., (Dung 1995; Caminada 2006a; Dung, Mancarella, and Toni 2007; Baroni, Giacomin, and Guida 2005)).

Axioms for Acceptability Semantics

A semantics aims at evaluating the overall strength (or acceptability) of every argument in an argumentation graph. This is done by aggregating the intrinsic strength of the argument and the overall strengths of its attackers. The overall strength of each attacker is based on its intrinsic strength and the overall strengths of its own attackers, an so on.

In what follows, we propose a set of 15 axioms that shed light on foundational concepts and principles behind semantics. The set is partitioned into three subsets of axioms: the first subset describes rationality properties of a semantics, the second one formalizes the *role* and the *impact* of attacks on the overall strength of an argument. The last subset of axioms describes key *factors* that may be taken into account when computing the overall strength.

Rationality Axioms

The first basic axiom ensures that the acceptability degree of an argument does not depend on its identity.

Axiom 1 (Anonymity) *A semantics* $\mathbf{S}$ *satisfies* anonymity *iff, for any two argumentation graphs* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *and* $\mathbf{A}' = \langle \mathcal{A}', \mathcal{R}' \rangle$*, for any isomorphism* f *from* $\mathbf{A}$ *to* $\mathbf{A}'$*, the following property holds:*

$$\forall\, a \in \mathcal{A}, \mathtt{Deg}^\mathbf{S}_\mathbf{A}(a) = \mathtt{Deg}^\mathbf{S}_{\mathbf{A}'}(f(a)).$$

This axiom is based on the smoothening assumption which ensures that nothing has changed when passing from graph $\mathbf{A}$ to graph $\mathbf{A}'$ (arguments have the same intrinsic strengths and attacks have the same weights).

The second basic axiom, called *independence*, states the following: the acceptability degree of an argument a should be independent of any argument or attack that is not connected to a (i.e., there is no path from that argument or attack to a, ignoring the direction of the edges).

Axiom 2 (Independence) *A semantics* $\mathbf{S}$ *satisfies* independence *iff, for any two argumentation graphs* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *and* $\mathbf{A}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ *such that* $\mathcal{A} \cap \mathcal{A}' = \emptyset$*, the following property holds:*

$$\forall\, a \in \mathcal{A},\ \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A} \oplus \mathbf{A}'}(a).$$

The next axiom, called *circumscription*, states that the acceptability degree of an argument should not depend on the arguments it itself attacks. This axiom does not depend on the smoothening assumption.

Axiom 3 (Circumscription) *A semantics* $\mathbf{S}$ *satisfies* circumscription *iff, for any two argumentation graphs* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *and* $\mathbf{A}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ *such that* $\mathcal{A} = \mathcal{A}'$ *and* $\mathcal{R}' = \mathcal{R} \cup \{(a, b)\}$ *with* $\mathtt{Att+}_{\mathbf{A}}(b) = \emptyset$*, the following holds: for all* $x \in \mathcal{A} \setminus \{b\}$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}'}(x)$.

It is worth noticing that Circumscription is a weaker version of the directionality axiom defined in (Baroni, Giacomin, and Guida 2005) for extension semantics.

The following axiom, called *monotony*, ensures that an argument cannot become stronger when its set of attackers gets bigger.

Axiom 4 (Monotony) *A semantics* $\mathbf{S}$ *satisfies* monotony *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if* $\mathtt{Att}_{\mathbf{A}}(a) \subseteq \mathtt{Att}_{\mathbf{A}}(b)$*, the following property holds:*

$$\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) \geq \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b).$$

This axiom is heavily based on the smoothening assumption. Indeed, since the two arguments a and b have the same intrinsic strengths, then their overall strengths may differ only if their attackers are different. Furthermore, since the role of an attack is to weaken its target, then the argument b may be weakened further by the attackers it does not share with the argument a.

The following mandatory axiom, called *symmetry*, states that arguments that are attacked by the same arguments receive the same degrees of acceptability.

Axiom 5 (Symmetry) *A semantics* $\mathbf{S}$ *satisfies* symmetry *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if* $\mathtt{Att}_{\mathbf{A}}(a) = \mathtt{Att}_{\mathbf{A}}(b)$*, the following property holds:*

$$\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b).$$

The next axiom, called *equivalence*, ensures that the overall strength of an argument depends solely on the overall strengths of its direct attackers. The overall strengths of the attackers are themselves evaluated on the basis of their direct attackers, and so on. Thus, the evaluation of an argument depends on the overall strengths of its direct and indirect attackers (respectively defenders).

Axiom 6 (Equivalence) *A semantics* $\mathbf{S}$ *satisfies* equivalence *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if there exists a bijective function* f *from* $\mathtt{Att}_{\mathbf{A}}(a)$ *to* $\mathtt{Att}_{\mathbf{A}}(b)$ *such that* $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$*, then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

This axiom holds under the smoothening assumption ensuring that a and b have the same intrinsic strength. Otherwise, even if the attackers of a and b have equal strengths, they may not have the same effect on both arguments.

Our last rationality axiom, called *neutrality*, gives a clear interpretation to value 0. It states that arguments that get this value have no impact on the arguments they attack. This axiom is also based on the smoothening assumption, namely the fact that arguments have the same intrinsic strength.

Axiom 7 (Neutrality) *A semantics* $\mathbf{S}$ *satisfies* neutrality *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if* $\mathtt{Att}_{\mathbf{A}}(b) = \mathtt{Att}_{\mathbf{A}}(a) \cup \{x\}$ *such that* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$*, then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Axioms on the Role and Impact of Attacks

The following axiom, called *maximality*, states that if an argument is not attacked, its overall strength is equal to its intrinsic strength (which is assumed equal to 1 in the paper).

Axiom 8 (Maximality) *A semantics* $\mathbf{S}$ *satisfies* maximality *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for any argument* $a \in \mathcal{A}$*, if* $\mathtt{Att}_{\mathbf{A}}(a) = \emptyset$*, then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = 1$.

The following axiom, called *weakening*, defines formally the role of attacks. It states that an attack weakens its target by decreasing its overall strength (possibly only by an infinitesimal amount). This is particularly true when the attack emanates from a non-rejected argument (i.e., an argument whose acceptability degree is greater than zero). This axiom is clearly based on the assumption that arguments have the same intrinsic strength.

Axiom 9 (Weakening) *A semantics* $\mathbf{S}$ *satisfies* weakening *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for any argument* $a \in \mathcal{A}$*, if* $\exists b \in \mathtt{Att}_{\mathbf{A}}(a)$ *such that* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) > 0$*, then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) < 1$.

The next axiom, called *weakening soundness*, states that the only way of decreasing the overall strength of an argument is by attacking the argument with a non-rejected argument. This axiom is based on the assumption that arguments have the same intrinsic strengths.

Axiom 10 (Weakening Soundness) *A semantics* $\mathbf{S}$ *satisfies* weakening soundness *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for any argument* $a \in \mathcal{A}$*, if* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) < 1$ *then* $\exists b \in \mathtt{Att}_{\mathbf{A}}(a)$ *such that* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) > 0$.

The previous axioms are about the role of attacks, which is weakening arguments. The next axiom concerns at what extent an attack may be harmful. We are particularly interested by the extreme case, i.e. whether or not an attack may reduce the acceptability degree of an argument to 0. We distinguish two opposite principles:

1. Attacks may lead to the rejection of arguments. This principle makes sense in some applications like defeasible reasoning. For instance, an argument built upon a default may be overruled by another argument which uses a more specific rule.
2. Arguments are *resilient* to attacks, and can never be completely rejected. This principle makes sense in practical applications like dialogue.

The next axiom, called *resilience*, captures the principle according to which an attack cannot reduce the acceptability degree of an argument to 0.

Axiom 11 (Resilience) *A semantics* $\mathbf{S}$ *satisfies* resilience *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for any argument* $a \in \mathcal{A}$, $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$.

This axiom separates the existing extension semantics from gradual (or ranking) semantics. Indeed, the former violate Resilience while the latter satisfy it. The choice of the semantics to use depends merely on the application at hand. This means there is no universal semantics. A semantics may be appropriate for a given application and not for another.

One may imagine several cases where an attack may reduce the acceptability degree of its target to 0. The aim of this paper is not to present an exhaustive list, but we provide one way which is already considered in the argumentation literature, namely by extension semantics. The axiom, called *killing*, says: any attack that comes from an argument with acceptability degree 1 leads to the rejection of its target.

Axiom 12 (Killing) *A semantics* $\mathbf{S}$ *satisfies* killing *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for any argument* $a \in \mathcal{A}$*, if* $\exists b \in \texttt{Att}_{\mathbf{A}}(a)$ *such that* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) = 1$*, then* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = 0$.

As we will see later, killing is the fundamental characteristic of extension semantics.

Axioms on Key Factors for Argument Evaluation

The following axioms introduce two key factors that may impact the overall strength of an argument: the *number* of non-rejected attackers of the argument and their *quality*. Recall that rejected attackers have no effect.

The more numerous the non-rejected attackers of an argument, the weaker the argument. We distinguish between two cases: the case where the argument has one non-rejected attacker and the case where it has several. This distinction allows a better understanding of the foundations of extension semantics. As we will see in a next section, these semantics are sensitive to the first case but not to the second one.

The first case is captured by an axiom called *triggering*. This axiom states that the overall strength of an argument must decrease once a first non-rejected attacker appears and there is room to decrease (i.e. the overall strength of the argument was higher than 0 before the introduction of the first non-rejected attacker). Indeed, if the argument was already rejected, the new attack cannot reject it to a greater extent.

Axiom 13 (Triggering) *A semantics* $\mathbf{S}$ *satisfies* triggering *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if*

- $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- $\forall x \in \texttt{Att}_{\mathbf{A}}(a)$, $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$,
- $\texttt{Att}_{\mathbf{A}}(b) = \texttt{Att}_{\mathbf{A}}(a) \cup \{y\}$ *and* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > 0$,

then $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

The second case is captured by an axiom called *counting*. It states that the more numerous the non-rejected attackers of an argument, the weaker the argument.

Axiom 14 (Counting) *A semantics* $\mathbf{S}$ *satisfies* counting *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$*, if*

- $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- $\exists x \in \texttt{Att}_{\mathbf{A}}(a)$ *such that* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) > 0$,
- $\texttt{Att}_{\mathbf{A}}(b) = \texttt{Att}_{\mathbf{A}}(a) \cup \{y\}$ *with* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > 0$,

then $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

The quality of non-rejected attackers is another factor that may impact the overall strength of an argument. The next axiom, called *reinforcement*, states that if the overall strength of an attacker is increased, then its target is weakened further provided that it is not already rejected.

Axiom 15 (Reinforcement) *A semantics* $\mathbf{S}$ *satisfies* reinforcement *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$*, if*

- $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- $\texttt{Att}_{\mathbf{A}}(a) \setminus \texttt{Att}_{\mathbf{A}}(b) = \{x\}$,
- $\texttt{Att}_{\mathbf{A}}(b) \setminus \texttt{Att}_{\mathbf{A}}(a) = \{y\}$,
- $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > \texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) > 0$,

then $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

One may wonder why $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x)$ cannot be 0. The reason is that case follows from Neutrality, Triggering and Counting.

Our last axiom, *boundedness*, is also about the quality of attacks. It states that if an argument is rejected, then it remains rejected if one of its attackers is strengthened.

Axiom 16 (Boundedness) *A semantics* $\mathbf{S}$ *satisfies* boundedness *iff, for any argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$*, for all* $a, b \in \mathcal{A}$ *such that*

- $\texttt{Att}_{\mathbf{A}}(a) \setminus \texttt{Att}_{\mathbf{A}}(b) = \{x\}$,
- $\texttt{Att}_{\mathbf{A}}(b) \setminus \texttt{Att}_{\mathbf{A}}(a) = \{y\}$,
- $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > \texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x)$

if $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = 0$*, then* $\texttt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) = 0$.

Naturally, the four previous axioms are based on smoothening assumption which ensures that the two arguments a and b have the same intrinsic strength.

Links and Compatibilities between Axioms

Each axiom introduces a novel concept or principle. There is almost no overlap between them. In other words, each axiom is primitive. Furthermore, they are all independent (none of them follows from another). There are nevertheless three notable exceptions. The first one concerns Maximality which follows from Weakening Soundness.

Proposition 1 *If a semantics $\mathbf{S}$ satisfies Weakening Soundness, then $\mathbf{S}$ satisfies Maximality.*

The converse is false. Indeed, some extension semantics satisfy Maximality while they violate Weakening Soundness.

The second exception concerns symmetry which follows from equivalence.

Proposition 2 *If a semantics $\mathbf{S}$ satisfies Equivalence, then $\mathbf{S}$ satisfies also Symmetry.*

The third exception concerns Monotony which follows from a subset of other axioms.

Proposition 3 *If a semantics $\mathbf{S}$ satisfies Independence, Directionality, Symmetry, Neutrality, Triggering and Counting, then $\mathbf{S}$ satisfies Monotony.*

Killing and Resilience are incompatible, i.e. there exists no semantics which can satisfy both axioms.

Proposition 4 *There exists no semantics which satisfies both Killing and Resilience.*

The remaining axioms are all compatible.

Proposition 5 *Anonymity, Independence, Circumscription, Monotony, Equivalence, Neutrality, Maximality, Weakening, Weakening Soundness, Resilience, Triggering, Counting, Reinforcement, and Boundedness are all compatible.*

Generalized Versions of Some Axioms

Neutrality, Triggering, Counting, Reinforcement and Boundedness are defined in a *basic* way. Indeed, the two arguments being compared are assumed to have the same attackers except one. The reason behind this choice of presentation is twofold: i) to have elementary axioms that are easy to grasp, and ii) the general version of each axiom follows from some axioms. By general version, we mean the case of arguments having arbitrary sets of attackers.

Let us start with the generalized version of Neutrality. The following result shows that it follows from Independence, Circumscription, Neutrality and Equivalence.

Proposition 6 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Neutrality and Equivalence, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, if there exists an injective function f from $\mathtt{Att}_{\mathbf{A}}(a)$ to $\mathtt{Att}_{\mathbf{A}}(b)$ s.t.*

- $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$, *and*
- $\forall y \in \mathtt{Att}_{\mathbf{A}}(b)$ *such that* $\nexists x \in \mathtt{Att}_{\mathbf{A}}(a)$ *with* $y = f(x)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) = 0$,

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Generalized Triggering follows from Independence, Circumscription, Monotony, Equivalence and Triggering.

Proposition 7 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Monotony, Equivalence, Triggering, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, if*

- $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$,
- $\mathtt{Att}_{\mathbf{A}}(b) = X \cup Y$ *such that*
  - $|X| = |\mathtt{Att}_{\mathbf{A}}(a)|$ *and* $\forall x \in X$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$,
  - $\forall y \in Y$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > 0$

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Generalized Counting follows from Independence, Circumscription, Monotony, Equivalence, and Counting.

Proposition 8 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Monotony, Equivalence, Counting, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, if*

- $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$ *and* $\exists x \in \mathtt{Att}_{\mathbf{A}}(a)$ *such that* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) > 0$
- $\mathtt{Att}_{\mathbf{A}}(b) = X \cup Y$ *such that*
  - *there exists a bijective function f from $\mathtt{Att}_{\mathbf{A}}(a)$ to X such that* $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,
  - $\forall y \in Y$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(y) > 0$

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Generalized version of Reinforcement follows from Independence, Circumscription, Boundedness, Equivalence, Reinforcement.

Proposition 9 *If semantics $\mathbf{S}$ satisfies Equivalence, Boundedness, Independence, Circumscription, Reinforcement, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, for all $a, b \in \mathcal{A}$, if*

- $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- $\mathtt{Att}_{\mathbf{A}}(a) = X \cup Z$,
- $\mathtt{Att}_{\mathbf{A}}(b) = Y \cup Z'$, *and*
- *there exists a bijective function f from X to Y such that* $\forall x \in X$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) < \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,
- *there exists a bijective function f' from Z to Z' such that* $\forall x \in Z$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Generalized Boundedness follows from Independence, Circumscription, Boundedness and Equivalence.

Proposition 10 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Boundedness and Equivalence, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$ such that*

- $\mathtt{Att}_{\mathbf{A}}(a) = X \cup Z$,
- $\mathtt{Att}_{\mathbf{A}}(b) = X' \cup Z'$,
- *there exists a bijective function f from X to X' such that* $\forall x \in X$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) < \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,
- *there exists a bijective function f' from Z to Z' such that* $\forall x \in Z$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,

if $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = 0$, *then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) = 0$.

Properties of Semantics Satisfying the Axioms

The aim of this section is to investigate properties of semantics that satisfy the axioms. We start by showing how key principles, on which extension semantics are based, can be decomposed into certain of our primitive axioms. In other words, we use our building blocks to reconstruct those principles shedding thus light on their foundations.

Extension semantics are based on a key principle, called *reinstatement*, according to which an argument can be accepted if its attackers are all rejected. This principle can be decomposed as follows:

Proposition 11 *Let $\mathbf{S}$ be a semantics which satisfies Independence, Circumscription, Neutrality, and Maximality. Let $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation graph. For any $a \in \mathcal{A}$ such that $\mathtt{Att}_{\mathbf{A}}(a) \neq \emptyset$, if $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = 1$.*

Another central notion of extension semantics is *defense*. Its basic idea is that the defenders (i.e., the attackers of the attackers) of an argument may improve the overall strength of the argument. The following result shows the foundations of defense. Indeed, if a semantics satisfies Independence, Circumscription, Maximality, Weakening, Boundedness, Equivalence and Reinforcement, then it considers defended arguments as stronger than non-defended ones.

Proposition 12 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Maximality, Weakening, Boundedness, Equivalence and Reinforcement, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, for all $a, b \in \mathcal{A}$, if*

- $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) < 1$,
- $|\mathtt{Att}_{\mathbf{A}}(a)| = |\mathtt{Att}_{\mathbf{A}}(b)|$,
- $\exists x \in \mathtt{Def}_{\mathbf{A}}(a)$ *s.t.* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) > 0$,
- $\mathtt{Def}_{\mathbf{A}}(b) = \emptyset$,

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

Note that in case $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b) = 1$, there is no room for a to be stronger than b.

Our next result shows a consequence of Anonymity axiom. It states that any semantics that satisfies Anonymity, assigns the same acceptability degree to all arguments of an elementary cycle. This shows that such semantics treat equally the arguments of elementary cycles.

Proposition 13 *If a semantics $\mathbf{S}$ satisfies Anonymity, then for every argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ such that $\mathbf{A}$ is an elementary cycle, the following property holds:*

$$\forall a, b \in \mathcal{A}, \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b).$$

A natural property that a semantics would satisfy is the so-called *Void Precedence* (VP) by Amgoud and Ben-Naim (2013). VP ensures that unattacked arguments are more acceptable than attacked ones. The next result shows the building blocks of VP.

Proposition 14 *If a semantics $\mathbf{S}$ satisfies Resilience, Maximality and Weakening, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, for all $a, b \in \mathcal{A}$, if*

- $\mathtt{Att}_{\mathbf{A}}(a) = \emptyset$,
- $\mathtt{Att}_{\mathbf{A}}(b) \neq \emptyset$,

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

We show next that if a semantics satisfies Independence, Circumscription, Monotony, Equivalence, Boundedness, and Reinforcement, then it also satisfies a nice property which says: if the attackers of argument b are at least as numerous and strong as those of argument a, then a is at least as strong as b. This property is the Counter-Transitivity (CT) postulate defined by Amgoud and Ben-Naim (2013). Our result shows thus its foundations.

Proposition 15 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Monotony, Equivalence, Boundedness, and Reinforcement, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, for all $a, b \in \mathcal{A}$, if there exists an injective function f from $\mathtt{Att}_{\mathbf{A}}(a)$ to $\mathtt{Att}_{\mathbf{A}}(b)$ such that $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) \leq \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) \geq \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.*

We also show that if the attackers of an argument b dominates the attackers of a both in terms of quality and quantity, then a is more acceptable than b.

Proposition 16 *If a semantics $\mathbf{S}$ satisfies Independence, Circumscription, Monotony, Equivalence, Boundedness and Reinforcement, then for any argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$, for all $a, b \in \mathcal{A}$, if*

- $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > 0$,
- *there exists an injective function $f : \mathtt{Att}_{\mathbf{A}}(a) \rightarrow \mathtt{Att}_{\mathbf{A}}(b)$ such that:*
  - $\forall x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) \leq \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$, *and*
  - $\exists x \in \mathtt{Att}_{\mathbf{A}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) < \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(f(x))$,

then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(b)$.

This section presented two kinds of results. First, it showed how our axioms capture crucial notions of extension semantics. Second, it presented some nice properties that semantics would enjoy if they satisfy the axioms.

Axiomatic Analysis of Extension Semantics

The aim of this section is to investigate the underpinnings of extension semantics, namely those proposed by Dung (1995). Before recalling the different semantics, let us first define the two basic concepts on which they are based.

Definition 6 (Conflict-freeness, Defence) *Let $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation graph and $\mathcal{E} \subseteq \mathcal{A}$.*

- *$\mathcal{E}$ is* conflict-free *iff $\nexists a, b \in \mathcal{E}$ such that $a\mathcal{R}b$.*
- *$\mathcal{E}$* defends *an argument a iff $\forall b \in \mathcal{A}$, if $b\mathcal{R}a$, then $\exists c \in \mathcal{E}$ such that $c\mathcal{R}b$.*

The following definition recalls the main semantics.

Definition 7 (Acceptability semantics) *Let $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation graph, and $\mathcal{E} \subseteq \mathcal{A}$ a conflict-free set.*

- *$\mathcal{E}$ is a* complete *extension iff it defends all its elements and contains any argument it defends.*
- *$\mathcal{E}$ is a* preferred *extension iff it is a maximal (w.r.t. set $\subseteq$) complete extension.*

- $\mathcal{E}$ *is a* stable *extension iff it attacks any argument in* $\mathcal{A} \setminus \mathcal{E}$.
- $\mathcal{E}$ *is a* grounded extension *iff it is a minimal (w.r.t. set* $\subseteq$*) complete extension.*

It is worth recalling that stable extensions may not exist. Furthermore, each stable extension is preferred, which is itself a complete extension.

Notations: $\mathtt{Ext}_x(\mathbf{A})$ denotes the set of all extensions of $\mathbf{A}$ under semantics x where $x \in \{p, s, c, g\}$ and p (respectively s, c, g) stands for preferred (respectively stable, complete, grounded). Since an argumentation graph $\mathbf{A}$ has a single grounded extension, we denote it by $\mathtt{GE}(\mathbf{A})$.

In the argumentation literature, the extensions of an argumentation graph are used for assigning an acceptability degree to each argument. The scale that is used is qualitative and contains three values: *sceptically accepted* (a degree which is assigned to arguments that belong to all extensions), *credulously accepted* (a degree which is assigned to arguments that belong to some but not all extensions), and *rejected* (a degree assigned to arguments that do not belong to any extension). This definition can be found in several papers like (Baroni and Giacomin 2007; Cayrol and Lagasquie-Schiex 2005; Grossi and Modgil 2015). In what follows, we will consider a more refined definition. The idea is to distinguish between two categories of arguments that do not belong to any extension: those that are not attacked by any extension, and those that are attacked by at least one extension. We will use thus a scale of 4 values $\{0, 0.3, 0.5, 1\}$. The value 1 refers to sceptically accepted arguments, 0.5 to credulously accepted arguments, 0.3 is assigned to arguments that do not belong to any extension and are not attacked by extensions, and 0 is assigned to rejected arguments that are attacked by at least one extension.

Definition 8 (Acceptability Degree) *Let* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *be an argumentation graph and* $x \in \{p, s, c, g\}$.
If $\mathtt{Ext}_x(\mathbf{A}) = \emptyset$*, then* $\forall a \in \mathcal{A}$, $\mathtt{Deg}^x_{\mathbf{A}}(a) = 0.3$. *Otherwise,*

- $\mathtt{Deg}^x_{\mathbf{A}}(a) = 1$ *iff for all* $\mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$, $a \in \mathcal{E}$.
- $\mathtt{Deg}^x_{\mathbf{A}}(a) = 0.5$ *iff* $\exists \mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$ *such that* $a \in \mathcal{E}$ *and* $\exists \mathcal{E}' \in \mathtt{Ext}_x(\mathbf{A})$ *such that* $a \notin \mathcal{E}'$.
- $\mathtt{Deg}^x_{\mathbf{A}}(a) = 0.3$ *iff for all* $\mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$, $a \notin \mathcal{E}$ *and* $\nexists \mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$ *such that* $\exists b \in \mathcal{E}$ *and* $b\mathcal{R}a$.
- $\mathtt{Deg}^x_{\mathbf{A}}(a) = 0$ *iff for all* $\mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$, $a \notin \mathcal{E}$ *and* $\exists \mathcal{E} \in \mathtt{Ext}_x(\mathbf{A})$ *such that* $\exists b \in \mathcal{E}$ *and* $b\mathcal{R}a$.

Since each argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ has a single grounded extension, then for all $a \in \mathcal{A}$, $\mathtt{Deg}^g_{\mathbf{A}}(a) \in \{0, 0.3, 1\}$. Furthermore, when the argumentation graph contains a finite number of arguments, the grounded extension is obtained by iterative application of a characteristic function to the empty-set as follows:

$$\mathtt{GE}(\mathbf{A}) = \bigcup_{i \geq 1} \mathcal{F}^i(\emptyset)$$

where for $X \subseteq \mathcal{A}$, $\mathcal{F}(X) = \{x \in \mathcal{A} \mid X \text{ defends } x\}$.

In what follows, we present a partial characterization of grounded semantics using our axioms. Indeed, we show that a semantics which satisfies Independence, Circumscription, Killing, Maximality and Neutrality, assigns value 1 to any argument belonging to the grounded extension and value 0 to any argument attacked by the grounded extension. However, nothing can be said about the remaining arguments.

Theorem 1 *Let* $\mathbf{S}$ *be a semantics which satisfies Independence, Circumscription, Killing, Maximality and Neutrality. For all argumentation graph* $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ *such that* $\mathcal{F}_{\mathbf{A}}(\emptyset) \neq \emptyset$*, the two following properties hold:*

- $\forall x \in \bigcup_{i \geq 1} \mathcal{F}^i_{\mathbf{A}}(\emptyset)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 1$,
- $\forall x \in \mathcal{A}$, *if* $\exists y \in \bigcup_{i \geq 1} \mathcal{F}^i_{\mathbf{A}}(\emptyset)$ *such that* $y\mathcal{R}x$*, then* $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{A}}(x) = 0$.

The next theorem shows for each of the recalled semantics the list of axioms it satisfies and the list of those it violates.

Theorem 2 *Table 1 summarizes the axioms that are satisfied (respectively violated) by grounded, stable, preferred, and complete semantics.*

From Table 1, the four semantics satisfy Anonymity, Monotony, Weakening, Defence Precedence and Killing. These axioms are at the heart of the family of extension semantics.

Unsurprisingly, since the four semantics satisfy Killing, they all violate Resilience. They also violate Void Precedence, Counter-Transitivity, and Counting. Consider the argumentation graph $\mathbf{A}_1$ depicted below:

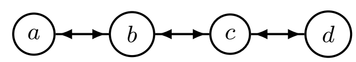

This graph has three stable extensions: $\mathcal{E}_1 = \{a, c\}$, $\mathcal{E}_2 = \{a, d\}$ and $\mathcal{E}_3 = \{b, d\}$. Thus, $\mathtt{Deg}^s_{\mathbf{A}_1}(a) = \mathtt{Deg}^s_{\mathbf{A}_1}(b) = \mathtt{Deg}^s_{\mathbf{A}_1}(c) = \mathtt{Deg}^s_{\mathbf{A}_1}(d) = 0.5$. However, $\mathtt{Att}_{\mathbf{A}_1}(c) = \{b, d\}$ while $\mathtt{Att}_{\mathbf{A}_1}(a) = \{b\}$.

Independence is satisfied by grounded, complete and preferred semantics. However, it is violated by stable semantics. The reason is the strong assumption which states that a stable extension should attack any argument left outside. This means that the evaluation of an argument may depend on arguments that are not related at all to the argument. Consider the argumentation graph $\mathbf{A}_2$ below.

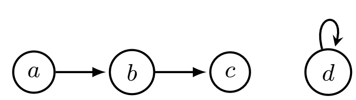

This graph has no stable extension and each of the 4 arguments gets degree 0.3. However, if we remove the loop, which is not connected at all to the other arguments, the remaining sub-graph has a stable extension $\{a, c\}$ and thus $\mathtt{Deg}^s_{\mathbf{A}_2}(a) = \mathtt{Deg}^s_{\mathbf{A}_2}(c) = 1$.

We believe that Independence is a mandatory property for ensuring precise evaluations. Assume that the argument d is about the weather in Toulouse, and the three other arguments (a, b, c) are about whether e-sport is a sport. Stable semantics mixes the evaluation of arguments which are about two different topics.

Another side-effect of the strong assumption behind stable semantics is the violation of Maximality. Indeed, when stable extensions do not exist, all the arguments get value 0.3 even non-attacked ones. Maximality is however satisfied by grounded, preferred and complete semantics.

Weakening Soundness is satisfied by grounded semantics and violated by the three other semantics. This means that stable (respectively preferred and complete) semantics does not evaluate an argument *solely* on the basis of the overall strengths of its attackers. An argument may be weakened even if it is not attacked or its attackers are all rejected. Thus, there is another factor at play in the evaluation of the arguments, namely *coalitions*. Each extension represents a coalition of arguments. According to those semantics, the overall strength of an argument represents whether or not the argument belongs to coalitions. Consider again the graph $\mathbf{A}_2$. Since it has no stable extension, then $\mathtt{Deg}^s_{\mathbf{A}_2}(a) = 0.3$, thus it is weakened even if it is not attacked at all.

Equivalence is another axiom which is satisfied only by grounded semantics. Under stable (respectively preferred and complete) semantic, two arguments may have different acceptability degrees even if their attackers have the same acceptability degrees. Consider the argumentation graph $\mathbf{A}_3$ depicted below:

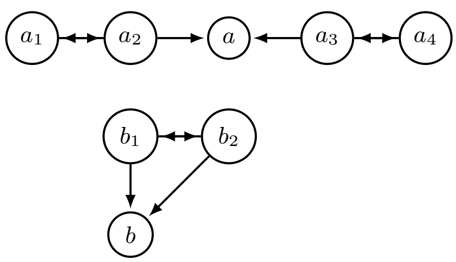

This graph has 8 stable extensions. It can be checked that $\mathtt{Deg}^s_{\mathbf{A}_3}(a_2) = \mathtt{Deg}^s_{\mathbf{A}_3}(a_3) = \mathtt{Deg}^s_{\mathbf{A}_3}(b_1) = \mathtt{Deg}^s_{\mathbf{A}_3}(b_2) = 0.5$. However, $\mathtt{Deg}^s_{\mathbf{A}_3}(a) = 0.5$ and $\mathtt{Deg}^s_{\mathbf{A}_3}(b) = 0$.

Remember that extension semantics are based on reinstatement principle according to which an argument may be accepted if all its attackers are rejected. This principle is used by Caminada (2006a) in his labeling functions. A labeling function assigns to each argument of an argumentation graph a label from the set $\{\mathtt{in}, \mathtt{out}, \mathtt{und}\}$. An argument is `in` if all its attackers are `out`, capturing thus reinstatement. Full correspondences have been shown between extensions (under the reviewed semantics) and different possible labellings of an argumentation graph.

At a first sight, one expects that Neutrality is satisfied by extension semantics since it says that rejected arguments have no effect on their targets. Surprisingly, this is not the case for preferred and complete semantics. The main reason is that those semantics may assign label `und` to arguments preventing thus the application of reinstatement. Since the labellings corresponding to stable extensions do never assign `und` to arguments (Caminada 2006a), Neutrality is satisfied by stable semantics. It is also satisfied by the grounded semantics since it ensures one extension.

Triggering counts the number of serious attackers only from 0 to 1. Indeed, it treats the case where an argument has only rejected attackers, then it receives a strong one. This axiom is satisfied by grounded and stable semantics but not by preferred and complete. This means that preferred and complete do not count at all. Consider the argumentation graph $\mathbf{A}_4$ depicted below.

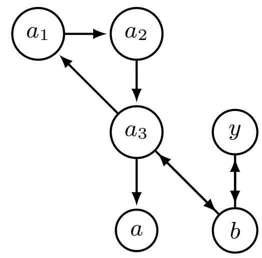

Clearly, $\mathtt{Att}_{\mathbf{A}_4}(b) = \mathtt{Att}_{\mathbf{A}_4}(a) \cup \{y\}$ with $\mathtt{Att}_{\mathbf{A}_4}(a) = \{a_3\}$. The graph $\mathbf{A}_4$ has two preferred extensions: $\mathcal{E}_1 = \{y\}$ and $\mathcal{E}_2 = \{a, b, a_1\}$. Thus, $\mathtt{Deg}^p_{\mathbf{A}_4}(a) = \mathtt{Deg}^p_{\mathbf{A}_4}(b) = 0.5$ while $\mathtt{Deg}^p_{\mathbf{A}_4}(y) = 0.5$ and $\mathtt{Deg}^p_{\mathbf{A}_4}(a_3) = 0$.

Table 1 shows the main differences between the evaluations returned by the four semantics. Stable and grounded semantics differ with respect to four axioms: Independence, Equivalence, Maximality and Weakening Soundness. Stable and preferred semantics are distinguished by Independence, Maximality, Neutrality, Triggering, Boundedness and Reinforcement. Indeed, stable semantics takes into account the quality of the attackers while preferred semantics neglects this factor. Stable semantics takes slightly the number of attackers into account (since it satisfies Triggering) while preferred semantics does not count at all. Finally, preferred and complete semantics satisfy the same set of axioms.

Axiomatic Analysis of h-categorizer Semantics

Gradual (or ranking) semantics are gaining increasing interest in the literature. Several such semantics were proposed (e.g. (Amgoud and Ben-Naim 2013; Matt and Toni 2008; Thimm 2012; Leite and Martins 2011; da Costa Pereira, Tettamanzi, and Villata 2011)). Such semantics do not compute extensions. They define mainly functions assigning a numerical value to each argument. This value represents the overall strength of an argument (i.e. its acceptability degree). Arguments are then ranked with regard to acceptability. Due to space limitation, we investigate the properties of only one such semantics: *h-categoriser* (Besnard and Hunter 2001). The latter assigns for every argument a of an argumentation graph $\mathbf{A} = \langle \mathcal{A}, \mathcal{R} \rangle$ an acceptability degree in the interval $(0, 1]$ as follows:

$$\mathtt{Deg}^h_{\mathbf{A}}(a) = \frac{1}{1 + \sum_{b \in \mathtt{Att}_{\mathbf{A}}(a)} \mathtt{Deg}^h_{\mathbf{A}}(b)}$$

with $\mathtt{Deg}^h_{\mathbf{A}}(a) = 1$ if $\mathtt{Att}_{\mathbf{A}}(a) = \emptyset$. This semantics is denoted by h.

Axioms - Semantics	Grounded	Stable	Preferred	Complete	h-Categorizer
Anonymity	•	•	•	•	•
Independence	•	○	•	•	•
Circumscription	•	○	•	•	•
Monotony	•	•	•	•	•
Equivalence	•	○	○	○	•
Neutrality	•	•	○	○	•
Maximality	•	○	•	•	•
Weakening	•	•	•	•	•
Weakening Soundness	•	○	○	○	•
Resilience	○	○	○	○	•
Killing	•	•	•	•	○
Triggering	•	•	○	○	•
Counting	○	○	○	○	•
Boundedness	•	•	○	○	•
Reinforcement	•	•	○	○	•
Void Precedence	○	○	○	○	•
Defence Precedence	•	•	•	•	•
Counter-Transitivity	○	○	○	○	•

Table 1: *Satisfaction of axioms by some existing semantics.* The symbol • (resp. ○) means the axiom is satisfied (resp. violated) by the semantics.

The overall strength of an argument depends on the overall strengths of its attackers, which themselves depend on the overall strengths of their own attackers, and so on. It is worth mentioning that this semantics was initially proposed for evaluating arguments of acyclic argumentation graphs. In (Pu et al. 2014), the authors extended the semantics to deal with any graph. Finally, the semantics is a particular case of the compensation-based semantics proposed in (Amgoud et al. 2016).

The following result shows the foundational ideas behind the evaluation made by h-categorizer semantics.

Theorem 3 *The last column of Table 1 summarizes the axioms that are satisfied (respectively violated) by h-categorizer semantics.*

h-categorizer semantics satisfies almost all the axioms except Killing. The reason of violating Killing is that arguments are resilient to attacks. Thus, an argument can never be be fully rejected (it cannot get degree 0). Consequently, reinstatement is not applicable. Note that Resilience is the fundamental axiom which separates this semantics from extension ones. It is also worth mentioning that Neutrality, Boundedness and Triggering are satisfied in a vacuous way since their conditions can never be satisfied (as 0 is not a possible acceptability degree). Finally, this semantics satisfies the properties of Void Precedence, Defence Precedence and Counter-Transitivity. In (Amgoud et al. 2016), it is shown that this semantics satisfies compensation. The idea is that a large number of weak attacks has the same effect as a smaller number of strong attacks.

The main axioms which separate h-categorizer semantics from Grounded semantics are Counting, Void Precedence ad Counter-Transitivity.

Conclusion

The contribution of this paper is fivefold: First, the paper introduced foundational concepts and principles of acceptability semantics (i.e. of the evaluation of arguments in an argumentation graph). It proposed a set of *primitive* axioms, i.e. axioms that cannot be further decomposed. Each axiom captures a precise idea, avoiding thus any overlapping between axioms. Another feature of the axioms is the fact that they are defined in *atomic* way, focusing thus on simple cases. This is particularly the case of Neutrality, Triggering, Counting, Reinforcement and Boundedness. The definitions are easy to grasp, and their general versions follow from the basic axioms.

The second contribution consists of investigating the properties of semantics that satisfy the axioms. In particular, we have shown that the key notions of reinstatement and defense follow from some of the axioms. This is of great importance since it shows the foundations of those notions.

The third contribution is a formal analysis of extension semantics, namely those proposed by Dung (1995). This analysis sheds light on the foundations of the semantics, and shows why they may return different evaluations. It is worth pointing out that comparative studies of the same semantics have been performed in the literature. However, what is compared is the extensions themselves and never the overall strengths of arguments under those semantics. This paper, provides to the best of our knowledge the first comparison of the overall strengths of arguments.

The fourth contribution is an axiomatic analysis of h-categorizer semantics. We have shown that it satisfies almost all the axioms except Killing.

The fifth contribution is a formal analysis of the difference between extension semantics and gradual semantics, namely h-categorizer semantics. We have shown that the main axioms separating the two families are Resilience and Counting. Extension semantics do not take into account the number of attackers, and an attack may lead to complete rejection of its target. This is not the case for h-categorizer semantics.

This work can be extended in several ways: First, we plan to fully characterize extension semantics using the axioms that are satisfied by those semantics and at least an additional axiom showing how arguments are killed (i.e., their degree is set to 0 after an attack). Killing axiom introduces one way, but as seen in the paper there is a second way which needs to be formalized. Another line of research consists of analyzing the other existing gradual semantics against the axioms. This will clarify the differences and similarities between them. A more ambitious goal would be to fully characterize the family of semantics that satisfy a given subset of axioms.

Acknowledgments

This work benefited from the support of CIMI and AMANDE ANR-13-BS02-0004 project funded by the French National Research Agency.

References

Amgoud, L., and Ben-Naim, J. 2013. Ranking-based semantics for argumentation frameworks. In *Proceedings of the 7th International Conference on Scalable Uncertainty Management, SUM'13*, 134–147.

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.

Amgoud, L.; Ben-Naim, J.; Doder, D.; and Vesic, S. 2016. Ranking arguments with compensation-based semantics. In *Proceedings of KR'16*.

Baroni, P., and Giacomin, M. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence* 171(10-15):675–700.

Baroni, P.; Giacomin, M.; and Guida, G. 2005. Scc-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence* 168:162–210.

Bench-Capon, T. 2003. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation* 13(3):429–448.

Besnard, P., and Hunter, A. 2001. A logic-based theory of deductive arguments. *Artificial Intelligence* 128(1-2):203–235.

Bonzon, E.; Maudet, N.; and Moretti, S. 2014. Coalitional games for abstract argumentation. In *Proceedings of Computational Models of Argument, COMMA'14*, 161–172.

Caminada, M. 2006a. On the issue of reinstatement in argumentation. In *Proceedings of the 10th European Conference on Logics in Artificial Intelligence JELIA'06*, 111–123.

Caminada, M. 2006b. Semi-stable semantics. In *Proceedings of the 1st International Conference on Computational Models of Argument, COMMA'06*, 121–130.

Cayrol, C., and Lagasquie-Schiex, M. 2005. Graduality in argumentation. *Journal of Artificial Intelligence Research* 23:245–297.

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

Dung, P.; Mancarella, P.; and Toni, F. 2007. Computing ideal skeptical argumentation. *Artificial Intelligence* 171:642–674.

Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Non-Monotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77:321–357.

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

Grossi, D., and Modgil, S. 2015. On the graded acceptability of arguments. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence, IJCAI'15*, 868–874.

Leite, J., and Martins, J. 2011. Social abstract argumentation. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI'11*, 2287–2292.

Matt, P., and Toni, F. 2008. A game-theoretic measure of argument strength for abstract argumentation. In *Proceedings of the 11th European Conference on Logics in Artificial Intelligence, JELIA'08*, 285–297.

Parsons, S.; Tang, Y.; Sklar, E.; McBurney, P.; and Cai, K. 2011. Argumentation-based reasoning in agents with varying degrees of trust. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS'11*, 879–886.

Pu, F.; Luo, J.; Zhang, Y.; and Luo, G. 2014. Argument ranking with categoriser function. In *Proceedings of the 7th International Knowledge Science, Engineering and Management Conference, KSEM'14*, 290–301.

Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *Proceedings of the 20th European Conference on Artificial Intelligence, ECAI'12*, 750–755.