https://doi.org/10.1016/j.ijar.2017.10.006

A Novel Weighted Defence and Its Relaxation in Abstract Argumentation

Stefano Bistarelli[a], Fabio Rossi[a], Francesco Santini[a,*]

[a]*Department of Mathematics and Computer Science, University of Perugia, Italy*

Abstract

When dealing with Weighted Abstract Argumentation, having weights on attacks clearly brings more information. The advantage, for instance, is the possibility to define a different notion of defence, checking also if the weight associated with it is stronger than the attack weight. In this work we study two different relaxations, one related to the new weighted defence we propose, by checking the difference between the composition of inward and outward attack-weights. The second one is related to how much inconsistency we are willing to tolerate inside an extension; such amount is computed by aggregating the costs of the attacks between any two arguments both inside an extension. These two relaxations are strictly linked: allowing a small conflict may lead to have more arguments into an extension, and consequently result in a stronger or weaker defence. Weights are represented by a semiring structure, which can be instantiated to different metrics used in the literature (*e.g.*, costs, probabilities, fuzzy levels).

Keywords: Abstract Argumentation Frameworks, Weighted attacks, Inconsistency tolerance, Relaxation, Semirings.

1. Introduction

An *Abstract Argumentation Framework* (AAF) [1] is represented by a pair $\langle \mathcal{A}_{rgs}, R \rangle$ consisting of a set of arguments and a binary relationship of attack defined among them. Given a framework, it is possible to examine the question on which set(s) of arguments can be accepted, hence collectively surviving the conflict defined by R. Answering this question corresponds to defining an argumentation semantics. The key idea behind *extension-based* semantics is to identify some sets of arguments (called *extensions*) that survive the conflict "together". A very simple example of AAF is $\langle \{a, b\}, \{R(a, b), R(b, a)\} \rangle$, where two arguments a and b attack each other. In this case, each of the two positions represented by either $\{a\}$ or $\{b\}$ can be intuitively valid, since no additional

*Corresponding author

Email addresses: `bista@dmi.unipg.it` (Stefano Bistarelli), `rossi@dmi.unipg.it` (Fabio Rossi), `francesco.santini@dmi.unipg.it` (Francesco Santini)

information is provided on which of the two attacks prevails. However, having weights on attacks results in such additional information, which can be fruitfully exploited in this direction. For instance, in case the attack $R(a, b)$ is stronger than (or preferred to) $R(b, a)$, taking the position defined by a may result in a better choice for an intelligent agent, since it can be defended better.

The aim of this work is to first *i)* propose a new notion of weighted defence and, from this, relax classically exact and sharp concepts in *Weighted Abstract Argumentation Frameworks* (*WAAFs*, see Section 11). This is accomplished by allowing *ii)* an internal conflict *inside* extensions satisfying a given extension-based semantics, and *iii)* by relaxing the defence of arguments w.r.t. the attacks coming from *outside* an extension. These are the three main results of the work.

The first goal is to provide a new definition of defence for WAAFs, here called *w-defence*, which encompasses weights in the style of similar works, as [2] and [3]. In our proposal, an extension $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ defends an argument $b \in \mathcal{A}_{rgs}$ from $a \in \mathcal{A}_{rgs}$, if the composition (a parametric $\otimes$ operation from a *c-semiring* structure [4]) of all the attack weights from $\mathcal{B}$ to a is stronger than the composition of *all* the attacks from a to $\mathcal{B}$. Differently from [2], where the arithmetic sum of all attack weights from $\mathcal{B}$ to a needs to be only stronger than the attack from a to b, we also consider the set of attacks from a to $\mathcal{B}$. Therefore, both our proposal and the one given by [2] suggest a collective defence from $\mathcal{B}$ to a, but, differently, in this paper we consider the group of attacks from a to $\mathcal{B}$ as a single entity, i.e., with a single global weight. We believe such a choice provides a more coherent view: in the literature, defence is usually checked by considering all the counter-attacks from a set $\mathcal{B}$ to a (*e.g.*, in order to satisfy admissibility), but each attack from a to $\mathcal{B}$ is treated separately (however, in case of fuzzy aggregation of weights, the two approaches are equivalent). Our intent is to normalise such dis-homogeneity.

Once having defined w-defence, we then proceed with the relaxation of the framework. Such two issues represent orthogonal relaxations, concerning either only internal arguments, or the relationship between the set of acceptable and not-acceptable arguments (interval versus external arguments). In this way, an autonomous reasoning-agent has more instruments to understand, for instance, whether tolerating a small conflict among its arguments considerably changes its point of view. As a possible scenario, a debate can be permeated by arguments advanced by *trolls* [5], which can accordingly generate noise in an abstract framework. Mitigating the inconsistency produced by them may let other arguments grow in strength (see Section 9). Internal inconsistency arises in many areas of AI and computing: merging information from heterogeneous sources, negotiation in multi-agent systems, or understanding natural language dialogues [6]. On the other hand, an agent could also be interested in defending its arguments with a higher or lower level of strength, in order to respectively strengthen the defence or increase the number of satisfied extensions. For instance, increasing both α and γ, something close to a stable extension can appear (there are AAF where the stable semantics is not satisfied).

In the end, we design α^{γ}-semantics (Section 6), where α is the amount of weight tolerated inside an extension satisfying it, and γ is the weight difference between a "full" defence (i.e., w-defence) and the "weaker" defence implemented by an α^γ-extension. We obtain them from classical formulations [1], hence we call them α^γ-conflict-free and α^γ-admissible sets, α^γ-complete, α^γ-preferred, and α^γ-stable semantics. These two parameters strictly influence each other: by relaxing α one can allow one or more new arguments be accepted into an extension, but, at the same time, their attacks contribute to the defence strength, i.e., to the computation of γ: if a is accepted in $\mathcal{B}$ because α is increased, $\mathcal{B} \cup \{a\}$ could be not γ-defensible anymore, or, on the other hand, it could become defensible even with a stricter γ. An agent can play with these two parameters with the purpose to "explore the neighbourhoods" of classical extensions, and take different decisions.

To have a general and formal representation of weights and operations on them (i.e., aggregation and preference), in this work we instead adopt a parametric algebraic framework based on c-semirings [4]. Hence, it is possible to consider different metrics within the same computational framework, as fuzzy or probabilistic scores, and model different kinds of AAFs in the literature (see Section 10). This represents a further result of the paper.

This paper elaborates on [7] and extends the works in [8] and [9]. All these papers are here integrated to offer an thorough view on the topic, by providing proofs, extended examples, and tests, which are not present in the single contributions; for instance, the case-study in Section 9.2 is new. The paper is structured as follows: Section 2.1 recollects the basic definitions of AAF given by [1], while in Section 2.2 we introduce c-semirings. Section 3 presents WAAFs, w-defence, and Section 4 reports a comparison with related notions in the literature [2, 3]. Section 5 relaxes w-defence by proposing γ-defence, where γ is the amount by which defence is weakened. In Section 6 we propose α^γ-semantics (*e.g.*, α^γ-stable), which extend classical ones by considering α and γ at the same time. Section 7 collects some formal results concerning such new semantics, *e.g.*, inclusion relations. In Section 8 we briefly describe an implementation of the proposed framework, together with some tests on random WAAFs. Section 9 presents two different case-studies to better motivate the formal results in the paper: one example is based on a real-world case-study, where the considered WAAF is directly extracted from *Amazon.com* reviews on a chosen product (a ballet tutu for kids). In Section 10 we describe related work, and, finally, Section 11 wraps up the paper by drawing final conclusions and suggesting future work.

2. Background

In the following of this section we first recollect the main definitions at the basis of AAFs [1] (Section 2.1), and then introduce c-semirings (Section 2.2). C-semirings represent a parametric framework where to deal with attack-weights. By changing the underlying c-semiring instantiation, it is possible to capture different metrics (*e.g.*, fuzzy or probabilistic ones).

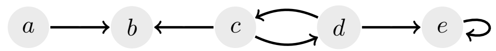

Figure 1: An example of AAF.

2.1. Abstract Argumentation Frameworks

In his pioneering work [1], Dung proposed *Abstract Frameworks* for Argumentation, where an argument is an abstract entity whose role is solely determined by its relations to other arguments:

Definition 1. *An Abstract Argumentation Framework (AAF) is a pair* $\langle \mathcal{A}_{rgs}, R \rangle$ *of a set* $\mathcal{A}_{rgs}$ *of arguments and a binary relation* R *on* $\mathcal{A}_{rgs}$*, called attack relation.* $\forall a_i, a_j \in \mathcal{A}_{rgs}$*,* $a_i R\, a_j$ *(or* $R(a_i, a_j)$*) means that* a_i *attacks* a_j *(*R *is asymmetric).*

An *argumentation semantics* is the formal definition of a method (either declarative or procedural) ruling the argument evaluation process. In the *extension*-based approach, a semantics definition specifies how to derive from an AAF a set of extensions, where an extension $\mathcal{B}$ of an AAF $\langle \mathcal{A}_{rgs}, R \rangle$ is simply a subset of $\mathcal{A}_{rgs}$. In Definition 2 we define conflict-free sets:

Definition 2 (Conflict-free sets)**.** *A set* $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ *is conflict-free iff no two arguments* a *and* b *in* $\mathcal{B}$ *exist such that* a *attacks* b*.*

All the following semantics rely (explicitly or implicitly) upon the concept of defence:

Definition 3 (Defence ($\mathbb{D}_0$))**.** *An argument* b *is defended by a set* $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ *(or* $\mathcal{B}$ *defends* b*) iff for any argument* $a \in \mathcal{A}_{rgs}$*, if* $R(a, b)$ *then* $\exists b \in \mathcal{B}$ *s.t.,* $R(b, a)$*.*

An admissible set of arguments [1] must be a conflict-free set that defends all its elements. Formally:

Definition 4 (Admissible sets)**.** *A conflict-free set* $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ *is admissible iff each argument in* $\mathcal{B}$ *is defended by* $\mathcal{B}$*.*

Three classical semantics [1] refining admissibility are defined in the following definitions:

Definition 5 (Complete semantics)**.** *An admissible set* $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ *is a complete extension iff each argument which is defended by* $\mathcal{B}$ *is in* $\mathcal{B}$*.*

Definition 6 (Preferred semantics)**.** *A preferred extension is a maximal (w.r.t. set inclusion) admissible subset of* $\mathcal{A}_{rgs}$*.*

Definition 7 (Stable semantics)**.** *A conflict-free set* $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ *is a stable extension iff for each argument which is not in* $\mathcal{B}$*, there exists an argument in* $\mathcal{B}$ *that attacks it.*

If $\sigma = \{cf, adm, com, prf, stb\}$ respectively stand for conflict-free and admissible sets, complete, preferred, and stable semantics, and $\xi_\sigma(F)$ is the set of extensions satisfying σ on a framework F, then we have $\xi_{stb}(F) \subseteq \xi_{prf}(F) \subseteq \xi_{com}(F) \subseteq \xi_{adm}(F) \subseteq \xi_{cf}(F)$. Moreover, for each σ except *stb* we have that $\xi_\sigma(F) \neq \emptyset$ holds.

The justification state of an argument can be conceived in terms of its extension membership: *justified* if it belongs to all the extensions satisfying a given semantics, *defensible* if it is in some extensions but not all of them, *overruled* if it does not belong to any of such extensions.

Definition 8 (Argument acceptance [10]). *Given any of the semantics $\sigma = \{adm, com, stb, prf\}$ and a framework F, an argument a is* i) justified *iff $\forall \mathcal{E} \in \sigma(F)$ (also $\mathcal{E}_{\sigma(F)}$), $a \in \mathcal{E}$,* ii) *a is defensible if $\exists \mathcal{E} \in \sigma(F), a \in \mathcal{E}$ and a is not justified,* iii) *a is overruled iff $\nexists \mathcal{E} \in \sigma(F), a \in \mathcal{E}$.*

Consider the AAF $F = \langle A, R\rangle$ in Figure 1, with $A = \{a, b, c, d, e\}$ and $R = \{(a, b), (c, b), (c, d), (d, c), (d, e), (e, e)\}$. We have that $adm(F)$ corresponds to $\{\emptyset, \{a\}, \{c\}, \{d\}, \{a, c\}, \{a, d\}\}$, $com(F) = \{\{a\}, \{a, c\}, \{a, d\}\}$, $prf(F) = \{\{a, c\}, \{a, d\}\}$, and $stb(F) = \{\{a, d\}\}$. Given such results, a is defensible in *adm(F)*, while it is justified in *com(F)*, *prf(F)*, and *stb(F)*.

2.2. C-semirings

In practice, c-semirings [4] are *commutative* ($\otimes$ is commutative) and *idempotent* semirings (i.e., $\oplus$ is idempotent), where $\oplus$ defines a complete lattice: every subset of elements have a *least upper bound*, or *lub*, and a *greatest lower bound*, or *glb*. In fact, c-semirings are semirings where $\oplus$ is used as a preference operator, while $\otimes$ is used to compose preference-values together.

Definition 9 (C-semirings [4]). *A commutative semiring is a tuple $\mathbb{S} = \langle S, \oplus, \otimes, \bot, \top\rangle$ such that S is a set, $\top, \bot \in S$, and $\oplus, \otimes : S \times S \to S$ are binary operators making the triples $\langle S, \oplus, \bot\rangle$ and $\langle S, \otimes, \top\rangle$ commutative monoids (semi-groups with identity), satisfying* i) *$\forall s, t, u \in S. s \otimes (t \oplus u) = (s \otimes t) \oplus (s \otimes u)$ (distributivity), and* ii) *$\forall s \in S. s \otimes \bot = \bot$ (annihilator). If $\forall s, t \in S. s \oplus (s \otimes t) = s$, the semiring is said to be absorptive. In short, c-semirings are defined as commutative and absorptive semirings.*

The idempotency of $\oplus$ leads to the definition of a partial ordering $\leqslant_\mathbb{S}$ over the set S (S is a poset). Such partial order is defined as $s \leqslant_\mathbb{S} t$ if and only if $s \oplus t = t$, and $\oplus$ returns the *lub* of s and t (defined also as $\sqcup$, while the *glb* is defined by $\sqcap$). This intuitively means that t is "better" than s. Some more properties can be derived on c-semirings [4]: *i)* both $\oplus$ and $\otimes$ are monotone over $\leqslant_\mathbb{S}$, *ii)* $\otimes$ is intensive (i.e., $s \otimes t \leqslant_\mathbb{S} s$), and *iii)* $\langle S, \leqslant_\mathbb{S}\rangle$ is a complete lattice. $\bot$ and $\top$ respectively are the bottom and top elements of such lattice. When also $\otimes$ is idempotent, *i)* $\oplus$ distributes over $\otimes$, *ii)* $\otimes$ returns the *glb* of two values in S, and *iii)* $\langle S, \leqslant_\mathbb{S}\rangle$ is a distributive lattice.

Well-known instances of c-semirings are: $\mathbb{S}_{boolean} = \langle\{false, true\}, \vee, \wedge, false, true\rangle$[1], $\mathbb{S}_{fuzzy} = \langle[0,1], \max, \min, 0,\ 1\rangle$, $\mathbb{S}_{bottleneck} = \langle\mathbb{R}^+ \cup \{+\infty\}, \max, \min, 0, \infty\rangle$, $\mathbb{S}_{probabilistic} = \langle[0,1],\ \max, \times, 0, 1\rangle$ (or *Viterbi* semiring), $\mathbb{S}_{weighted} = \langle\mathbb{R}^+ \cup \{+\infty\}, min, +, +\infty, 0\rangle$.

Furthermore, it is also possible to consider several optimisation criteria at the same time: the Cartesian product of semirings is still a semiring. Clearly, in this case the ordering induced by $\oplus$ is partial. For instance, the two couples $\langle s_1, s_2\rangle$ and $\langle s_3, s_4\rangle$ with $s_1, s_3 \in S_1$, $s_2, s_4 \in S_2$, and $s_1 \leqslant_{\mathbb{S}_1} s_3$ while $s_2 \geqslant_{\mathbb{S}_2} s_4$, are not directly comparable.

Although c-semirings have been historically used as monotonic structures where to aggregate costs (and find best solutions), the need of removing values has raised in local consistency algorithms and non-monotonic algebras using constraints (*e.g.*, [11]). A solution comes from *residuation theory* [12], a standard tool on tropical arithmetics that allows for obtaining a division operator that represents a "weak" inverse of $\otimes$.

Definition 10 (Residuation [11])**.** *Let $\mathbb{S}$ be a tropical semiring. $\mathbb{S}$ is residuated if the set $\{x \in S \mid t \otimes x \leqslant_{\mathbb{S}} s\}$ admits a maximum for all elements $s, t \in S$, such that $s \leqslant_{\mathbb{S}} t$.*

The maximal element among solutions is denoted denoted as $s \oslash t$. Since a complete[2] tropical-semiring is also residuated, we have that all the classical instances of c-semiring presented above are residuated. A semiring is invertible if there exists an element $r \in S$ such that $t \otimes r = s$ for all elements $s, t \in A$ such that $s \leqslant_{\mathbb{S}} t$. $\mathbb{S}$ is uniquely invertible if r is unique whenever $t \neq \perp$.

Definition 11 (Unique invertibility [11])**.** *If $\mathbb{S}$ is absorptive and invertible, then it is uniquely invertible iff it is cancellative, i.e., $\forall s, t, u \in S.(s \otimes u = t \otimes u) \wedge (u \neq 0) \Rightarrow s = t$.*

Since all the previously listed instances of c-semirings are cancellative, they are uniquely invertible as well. For instance,

$$\min\{x \mid b + x \geqslant a\} = \begin{cases} 0 & \text{if } b \geqslant a \\ a - b & \text{if } a > b \end{cases} \qquad \textit{weighted}$$

$$\max\{x \mid \min(b, x) \leqslant a\} = \begin{cases} 1 & \text{if } b \leqslant a \\ a & \text{if } a < b \end{cases} \qquad \textit{fuzzy}$$

3. Weighted Abstract AFs

In this section we rephrase some of the classical definitions given in Section 2.1, with the purpose to parametrise them with the notion of weighted at-

[1] Boolean c-semirings can be used to model crisp problems and classical Argumentation [1].

[2] $\mathbb{S}$ is complete if it is closed with respect to infinite sums, and the distributivity law holds also for an infinite number of summands [11].

tack and c-semiring. Such notions, *e.g.*, the one of w-defence, are the premises behind the new semantics we then propose in Section 6. The following definition reshapes the notion of WAAF into *semiring-based WAAF* [7], called $WAAF_{\mathbb{S}}$:

Definition 12 (Semiring-based WAAF [7])**.** *A semiring-based Weighted AAF ($WAAF_{\mathbb{S}}$) is a quadruple $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, where $\mathbb{S}$ is a c-semiring $\langle S, \oplus, \otimes, \bot, \top\rangle$, $\mathcal{A}_{rgs}$ is a set of arguments, R the attack binary-relation on $\mathcal{A}_{rgs}$, and $W : \mathcal{A}_{rgs} \times \mathcal{A}_{rgs} \longrightarrow S$ is a binary function. Given $a, b \in \mathcal{A}_{rgs}$ and $R(a, b)$, then $W(a, b) = s$ means that a attacks b with a weight $s \in S$. Moreover, we require that $R(a, b)$ iff $W(a, b) <_{\mathbb{S}} \top$.*

In Figure 2 we provide an example of a weighted AAF describing the $WAAF_{\mathbb{S}}$ defined by $\mathcal{A}_{rgs} = \{a, b, c, d, e\}$, $R = \{(a, b), (c, b), (c, d), (d, c), (d, e), (e, e)\}$, with $W(a, b) = 7, W(c, b) = 8, W(c, d) = 9, W(d, c) = 8, W(d, e) = 5, W(e, e) = 6$, and $\mathbb{S} = \langle \mathbb{R}^+ \cup \{\infty\}, \min, +, \infty, 0\rangle$ (i.e., the weighted semiring).

Therefore, each attack is associated with a semiring value that represents the "strength" of an attack between two arguments. We can consider the weights in Figure 2 as supports to the associated attack, as similarly suggested in [13]. A semiring value equal to the top element of the c-semiring $\top$ (*e.g.*, 0 for the weighted semiring) represents a no-attack relation between two arguments: for instance, $(a, c) \notin R$ in Figure 2 corresponds to $W(a, c) = 0$. Note that, whenever there is an attack between two arguments, its weight is different from $\top$: for example, $W(a, b) = 7$ in Figure 2. On the other side, the bottom element, i.e., $\bot$ (*e.g.*, ∞ for the weighted semiring), represents the strongest attack possible.

In Definition 13 we define the attack strength for a set of arguments that attacks an argument, a different set of arguments, or an argument that attacks a set of arguments; the former and the latter are what we need to define w-defence. In the following, we will use $\bigotimes$ to indicate the $\otimes$ operator of the c-semiring $\mathbb{S}$ on a set of values:

Definition 13 (Attacks to/from sets of arguments)**.** *Given a $WAAF_{\mathbb{S}}$, $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$,*

- *a set of arguments $\mathcal{B}$ attacks an argument a with a weight of $k \in S$ if*

$$W(\mathcal{B}, a) = \bigotimes_{b \in \mathcal{B}} W(b, a) = k$$

- *an argument a attacks a set of arguments $\mathcal{B}$ with a weight of $k \in S$ if*

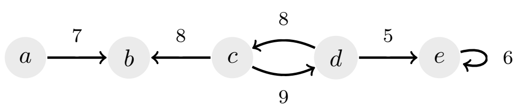

Figure 2: An example of WAAF, adding weights to Figure 1.

$$W(a, \mathcal{B}) = \bigotimes_{b \in \mathcal{B}} W(a, b) = k$$

- *a set of arguments $\mathcal{B}$ attacks a set of arguments $\mathcal{D}$ with a weight $k \in S$ if*

$$W(\mathcal{B}, \mathcal{D}) = \bigotimes_{b \in \mathcal{B}, d \in \mathcal{D}} W(b, d) = k$$

For example, looking at Figure 2 we have that $W(\{a, c\}, b) = 15$, $W(c, \{b, d\}) = 17$, and $W(\{a, c\}, \{b, d\}) = 24$.

We are now ready to define our version of weighted defence, i.e., w-defence:

Definition 14 (w-defence ($\mathbb{D}_w$)). *Given a $WAAF_{\mathbb{S}}$, $WF = \langle \mathcal{A}_{rgs}, R, W, S \rangle$, $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ w-defends $b \in \mathcal{A}_{rgs}$ iff $\forall a \in \mathcal{A}_{rgs}$ such that $R(a, b)$, we have that*

$$W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$$

As previously advanced, a set $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ w-defends an argument b from a, if the $\otimes$ of all attack weights from $\mathcal{B}$ to a is worse[3] (w.r.t. $\leqslant_{\mathbb{S}}$) than the $\otimes$ of the attacks from a to $\mathcal{B} \cup \{b\}$.

For example, the set $\{c\}$ in Figure 2 defends c from d because $W(d, \{c\}) \geqslant_{\mathbb{S}} W(\{c\}, d)$, i.e., $(8 \leqslant 9)$. On the other hand, $\{d\}$ in Figure 2 does not w-defend d because $W(c, \{d\}) \ngeqslant_{\mathbb{S}} W(\{d\}, c)$.

As defined, w-defence implies the classical Dung's defence in Definition 3:

Proposition 1 ($\mathbb{D}_w \Rightarrow \mathbb{D}_0$). *Given a $WAAF_{\mathbb{S}}$, $WF = \langle \mathcal{A}_{rgs}, R, W, S \rangle$, a subset of arguments $\mathcal{B}$, and $b \in \mathcal{A}_{rgs}$, "$\mathcal{B}$ w-defends b" (Definition 14) $\Rightarrow$ "$\mathcal{B}$ defends b" (Definition 3) in the corresponding not-weighted $\langle \mathcal{A}_{rgs}, R \rangle$.*

Proof. As hypothesis we have $R(a, b)$ (from Definition 14), then $W(a, \mathcal{B} \cup \{b\}) \neq \top$. Therefore, if $W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$ is true (i.e., $\mathcal{B}$ w-defends b from a), this implies that $W(\mathcal{B}, a) \neq \top$. This can be also read as "$\mathcal{B}$ attacks a", which exactly corresponds to the original definition of defence (see Definition 3). $\square$

Moreover, the following proposition equates defence and w-defence in case we adopt the boolean c-semiring (see Section 2.2):

Proposition 2. *Given a $WAAF_{\mathfrak{S}}$, $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S} \rangle$, where $\mathbb{S} = \langle \{true, false\}, \vee, \wedge, false, true \rangle$ (i.e., the boolean semiring), "$\mathcal{B}$ w-defends a" $\Longleftrightarrow$ "$\mathcal{B}$ defends a".*

Proof. This holds because, $\mathcal{B}$ defends b corresponds to, "if $W(a, b) \neq \top$ then $W(\mathcal{B}, a) \neq \top$". But, since we are using the boolean semiring, this statement can only correspond to, "if $W(a, b) =$ *false* then $W(\mathcal{B}, a) =$ *false*", since the set of

[3] Note that, when considering the partial order of a generic semiring, we will often use "worse" or "better" because "greater" or "lesser" would be misleading: in the weighted semiring, $7 \leqslant_{\mathbb{S}} 3$, i.e., lesser means better.

preferences only contains $\top$ (*true*) and $\perp$ (*false*). Therefore, $W(a, b) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$ is always true (in this case, *false* $\geqslant_{\mathbb{S}}$ *false*), and $\mathcal{B}$ w-defends b from a. $\square$

4. Two Related Definitions

Two of the most related definitions of weighted defence (i.e., Definition 14) are presented in [2] and [3]. In the following we condense their main features and we show how our approach differs.

In [3] attacks are relatively ordered by their force, i.e., $R(a, b) \gg R(b, a)$ means that the former attack is stronger than the latter (vice-versa, a weaker attack). Equivalent and incomparable classes are considered as well, i.e., $R(a, b) \approx R(b, a)$ and $R(a, b)?R(b, a)$, respectively. This is accordingly reflected by the defence definition, where considering $R(a, b)$ and $R(c, a)$ we can have that c is a *strong*, *weak*, *normal*, or *unqualified* defender of b. Therefore, an argument b is defended by $\mathcal{B}$ if, and only if, for any argument a such that $R(a, b)$, there is an argument $c \in \mathcal{B}$ such that $R(c, a)$, and according to the desired defence strength, $R(c, a) \gg R(a, b)$, $R(c, a) \ll R(a, b)$, $R(c, a) \approx R(a, b)$, and $R(c, a)?R(a, b)$. For instance, when requiring a level $[\gg, \approx]$, for each attacker a of b there must must be either a strong or a normal defender $c \in \mathcal{B}$. In Definition 15 we exactly rephrase such defence by modelling the total order defined by $[\gg, \approx]$ with a c-semiring $\mathbb{S}$:

Definition 15 (Defence $\mathbb{D}_1$ [3])**.** *Given $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S} \rangle$, $a, b, c \in \mathcal{A}_{rgs}$, $\mathcal{B} \subseteq \mathcal{A}_{rgs}$, then b is defended by $\mathcal{B}$ if $\forall R(a, b), \exists c \in \mathcal{B}$ s.t. $W(a, b) \geqslant_{\mathbb{S}} W(c, a)$.*

In [2] the authors define $\sigma^{\boxtimes}$-extensions, where σ is one of the given semantics (*e.g.*, admissible), and $\boxtimes$ is an *aggregation function* ($\otimes$ in a c-semiring). $\boxtimes$ needs to satisfy non-decreasingness, minimality, and identity:[4] two examples in the paper are the arithmetic sum and max. Even the notion of defence is refined: in Definition 16 we cast it in the same semiring-based framework.

Definition 16 (Defence $\mathbb{D}_2$ [2])**.** *Given $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S} \rangle$, an argument b is defended by a subset of arguments $\mathcal{B}$ if $\forall a \in \mathcal{A}_{rgs}$ s.t. $R(a, b)$, we have that $W(a, b) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$.*

Thus, an argument b is $\boxtimes$-acceptable if for each attack from a against b, the aggregated weight of the collective defence of b is greater than $W(a, b)$, according to $\mathbb{D}_2$. Such phrasing of defence is also equivalent to what presented in [7].

By using the same semiring-based framework, it is now possible to relate such notions of defence together (we remind that $\mathbb{D}_w$ stands for w-defence).

Theorem 4.1. $\mathbb{D}_w \Rightarrow \mathbb{D}_2$.

Proof. If $W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$ (i.e., Definition 14 holds) then $W(a, b) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$ (also Definition 16 holds), due to $W(a, b) \geqslant_{\mathbb{S}} W(a, \mathcal{B} \cup \{b\})$, because of the monotonicity of $\otimes$ operator (see Section 2.2). $\square$

[4] Such properties are satisfied by a c-semiring (see Section 2.2).

Moreover, we can link $\mathbb{D}_1$ and $\mathbb{D}_2$ as well:

Theorem 4.2. $\mathbb{D}_1 \Rightarrow \mathbb{D}_2$.

Proof. This is equivalent to prove that $\forall R(a,b), \exists c \in \mathcal{B}$ s.t. $W(a,b) \geqslant_{\mathbb{S}} W(c,a) \Rightarrow W(a,b) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$. If such c exists, we also know that $W(c,a) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$ holds, given Definition 13 and monotonicity of $\otimes$; transitivity leads to $W(a,b) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$, proving $\Rightarrow$. □

In case the c-semiring we use is the fuzzy one, i.e., $\langle [0,1], \max, \min, 0, 1\rangle$, then $\mathbb{D}_2$ collapses into $\mathbb{D}_1$, as Theorem 4.3 states.

Theorem 4.3. *If* $\mathbb{S} = \langle [0,1], \max, \min, 0, 1\rangle$*, then* $\mathbb{D}_1 \Leftrightarrow \mathbb{D}_2$.

Proof. This is equivalent to prove that $\forall R(a,b), \exists c \in \mathcal{B}$ s.t. $W(a,b) \geqslant_{\mathbb{S}} W(c,a) \Leftrightarrow W(a,b) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$. $\Rightarrow$ can be proved for any semiring $\mathbb{S}$ (see Theorem 4.2). In order to prove $\Leftarrow$, $W(\mathcal{B},a)$ is computed by using Definition 13 and min, hence there exists at least one $c \in \mathcal{B}$ s.t. $W(a,b) \geqslant_{\mathbb{S}} W(c,a)$. □

This result permits to relate $\mathbb{D}_w$ and $\mathbb{D}_1$ when using the fuzzy c-semiring:

Corollary 4.4. *If* $\mathbb{S} = \langle [0,1], \max, \min, 0, 1\rangle$*, then* $\mathbb{D}_w \Rightarrow \mathbb{D}_1$.

Proof. This directly follows from Theorem 4.1 when using a fuzzy c-semiring, i.e., Theorem 4.3. □

To conclude, we show that all the three $\mathbb{D}_w$, $\mathbb{D}_1$, and $\mathbb{D}_2$ collapse to the classical defence $\mathbb{D}_0$ [1] when considering the framework without weights.

Theorem 4.5. *If* $\mathbb{S} = \langle \{true, false\}, \vee, \wedge, false, true\rangle$*, then* $\mathbb{D}_w \Leftrightarrow \mathbb{D}_0 \Leftrightarrow \mathbb{D}_1 \Leftrightarrow \mathbb{D}_2$.

Proof. $\mathbb{D}_w \Leftrightarrow \mathbb{D}_0$ is proved in Proposition 2. To show $\mathbb{D}_w \Leftrightarrow \mathbb{D}_2$ we only need to prove $\mathbb{D}_2 \Rightarrow \mathbb{D}_w$ ($\Leftarrow$ holds from Theorem 4.1): this holds because if $W(a,b) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$ it means that if $W(a,b)$ is *false* then $W(\mathcal{B},a)$ is *false* (due to $\geqslant_{\mathbb{S}}$); hence, $W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$, i.e., $\mathbb{D}_w$, holds as well. To show $\mathbb{D}_1 \Leftrightarrow \mathbb{D}_2$ we only need to prove $\mathbb{D}_2 \Rightarrow \mathbb{D}_1$ ($\Leftarrow$ holds from Theorem 4.2): similarly, if $W(\mathcal{B},a)$ is *false*, than $\exists c \in \mathcal{B}$ s.t. $W(a,b) \geqslant_{\mathbb{S}} W(c,a)$, since $\exists c \in \mathcal{B}$ s.t. $W(c,a) = false$ (i.e., c attacks a). □

An example on how $\mathbb{D}_w$, $\mathbb{D}_1$, and $\mathbb{D}_2$ differently work is provided in Figure 3. We read this example by considering the weighted c-semiring, i.e., $\mathbb{S} = \langle \mathbb{R}^+ \cup \{+\infty\}, min, +, +\infty, 0\rangle$. Argument b is defended by $\mathcal{B} = \{b, c, d, e\}$ according to $\mathbb{D}_w$ and $\mathbb{D}_2$ (consequently, according to Theorem 4.2), since $W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B},a)$ ($0.7 \leqslant 0.8$). It is not defended according to $\mathbb{D}_1$, since $W(d,a) \geqslant_{\mathbb{S}} W(a,b)$ ($0.6 \leqslant 0.7$) and $W(c,a) \geqslant_{\mathbb{S}} W(a,b)$ ($0.2 \leqslant 0.7$).

On the other hand, considering the attacks from f to $\mathcal{B}$ instead, $\mathbb{D}_w$ does not hold: $W(f,d) \otimes W(f,e) \ngeqslant_{\mathbb{S}} W(d,f) \otimes W(c,f)$ (i.e., $0.8 \nleqslant 0.7$). However, $\mathbb{D}_2$ holds because $W(f,d) \geqslant_{\mathbb{S}} W(\{d,c\},f)$ (i.e., $0.5 \leqslant 0.7$) and $W(f,e) \geqslant_{\mathbb{S}}$

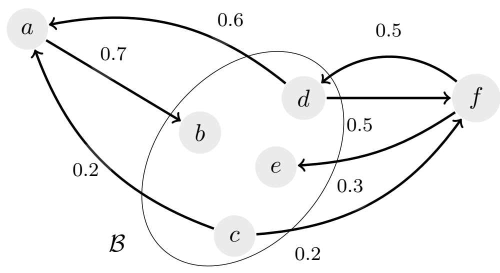

Figure 3: An example of WAAF where $\{b, c, d, e\}$ are defended by $\mathcal{B}$ according to $\mathbb{D}_2$ (using the weighted semiring).

$W(\{d, c\}, f)$ (i.e., $0.3 \leqslant 0.7$). Therefore, considering the WAAF in Figure 3, $\mathcal{B}$ defends itself from a and f according to only $\mathbb{D}_2$ ($\mathbb{D}_w$ and $\mathbb{D}_1$ do not hold).

Reading the same example in Figure 3 with $\mathbb{S} = \langle [0,1], \max, \min, 0, 1\rangle$ instead, $\mathbb{D}_2$ collapses to $\mathbb{D}_1$ (Theorem 4.3) and $\mathbb{D}_1$ does not hold due to $R(a, b)$. According to Theorem 4.4, since $\mathbb{D}_1$ is not valid then $\mathbb{D}_w$ cannot hold as well.

This example introduces us to relating admissible sets (at the core of all the extensions in [1]) using $\mathbb{D}_w$, $\mathbb{D}_1$, and $\mathbb{D}_2$.

Admissibility with $\mathbb{D}_w$*,* $\mathbb{D}_1$*,* $\mathbb{D}_2$*, and* $\mathbb{D}_0$. In the following we study how admissible sets are related considering $\mathbb{D}_w$, $\mathbb{D}_1$, and $\mathbb{D}_2$. We focus on this semantics because it is at the core of the other ones proposed by [1] (see Section 2.1), explicitly (i.e., complete, preferred), or implicitly (i.e., stable). We respectively call adm_1 and adm_2 the set of admissible sets using $\mathbb{D}_1$ and $\mathbb{D}_2$, adm_w is our proposal (Section 6), and adm_0 is the classical definition [1].

Theorem 4.6. *Given* $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$ *where* $\mathbb{S} = \langle [0,1], \max, \min, 0, 1\rangle$*, then* $adm_w(WF) = adm_1(WF) = adm_2(WF) \subseteq adm_0(WF)$*.*

Proof. $adm_1(WF) = adm_2(WF)$ directly derives from $\mathbb{D}_1 \Leftrightarrow \mathbb{D}_2$ (see Theorem 4.3). In order to prove $adm_w(WF) = adm_1(WF) = adm_2(WF)$, since we have already proved Theorem 4.4, we only need to show that $\mathbb{D}_1 \Rightarrow \mathbb{D}_w$, i.e., $\forall R(a, b), \exists c \in \mathcal{B}$ s.t. $W(a, b) \geqslant_{\mathbb{S}} W(c, a) \Rightarrow W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$. Therefore, $\exists c \in \mathcal{B}$ s.t. $W(a, b) \geqslant_{\mathbb{S}} W(a, \mathcal{B}) \geqslant_{\mathbb{S}} W(c, a)$, since $W(a, \mathcal{B})$ is the worst (min) of the attacks from a to $\mathcal{B}$. Given $W(c, a) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$ (see Corollary 4.4), consequently we have $W(a, \mathcal{B}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$. Finally, due to Proposition 1, we have the last inclusion of the theorem, i.e., $adm_w(WF) \subseteq adm_0(WF)$. □

5. Relaxing w-defence

The notion of w-defence given in Definition 14 can be relaxed in order to meet the one in [2] (see Definition 16) and, ultimately, the classical defence given by [1]. The result is here called γ-defence, and it is parametrised on a threshold-value γ: such γ is used to consider arguments that are not "fully" w-defended, i.e., for which $W(a, \mathcal{B} \cup \{b\}) \not\geqslant_{\mathbb{S}} W(\mathcal{B}, a)$:

Definition 17 (γ-defence ($\mathbb{D}_\gamma$))**.** *Given $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S} = \langle S, \oplus, \otimes, \bot, \top \rangle \rangle$ and $\gamma \in S$, $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ γ-defends $b \in \mathcal{A}_{rgs}$ iff $\forall a \in \mathcal{A}_{rgs}$ such that $R(a, b)$ we have that $W(\mathcal{B}, a) \neq \top$ and*

$$\Big(W(a, \mathcal{B} \cup \{b\}) \oslash W(\mathcal{B}, a)\Big) \geqslant_{\mathbb{S}} \gamma$$

Considering the example in Figure 3, for instance $\mathcal{B}$ 1-defends d from f (i.e., $\gamma = 1$):

$$\Big(W(a, \mathcal{B} \cup \{b\}) - W(\mathcal{B}, a)\Big) \leqslant 1,$$

since $8 - 7 = 1$ and $1 \leqslant 1$.

Remark 5.1. *Note that, not considering the comparison against γ, Definition 17 is identical to Definition 14 except for an additional condition, i.e., $W(\mathcal{B}, a) \neq \top$, which means that $\mathcal{B}$ has to attack a, i.e., $\exists b \in \mathcal{B}$ s.t. $R(b, a)$. This attack is automatically implied in Definition 14, but it has to be explicitly stated in Definition 14. Otherwise, $\mathcal{B}_1 = \{b\}$ and $\mathcal{B}_2 = \{d\}$ in Figure 4 would be treated in the same way and 1-defend b and d respectively. $\mathcal{B}_2$ has not to defend d instead, since there is not any attack from $\mathcal{B}_2$ to c. We opt for this choice because we want to reconnect to the classical defence of Dung ($\mathbb{D}_0$ [1]) when $\gamma = \bot$ (see Proposition 4).*

Next proposition links γ-defence to w-defence.

Proposition 3. *$\top$-defence ($\gamma = \top$) is equivalent to w-defence, i.e., $\mathbb{D}_\top \Leftrightarrow \mathbb{D}_w$.*

Proof. $\mathcal{B}$ $\top$-defends b if $W(a, \mathcal{B} \cup \{b\}) \oslash W(\mathcal{B}, a) \geqslant_{\mathbb{S}} \top$, which is true iff $W(a, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} W(\mathcal{B}, a)$, which corresponds to the definition of w-defence (see Definition 14). □

Next proposition shows how it is possible to reconnect γ-defence to Dung's original definition of defence.

Proposition 4. *$\bot$-defence ($\gamma = \bot$) is equivalent to the original definition of defence given by Dung [1], i.e., $\mathbb{D}_\bot \Leftrightarrow \mathbb{D}_0$.*

Proof. In γ-defence (see Definition 17), $\mathcal{B}$ has to defend all its arguments as for [1]. Moreover, using $\gamma = \bot$ allows us to forget weights of attacks and counter-attacks, since they can freely range from $\top$ to $\bot$. □

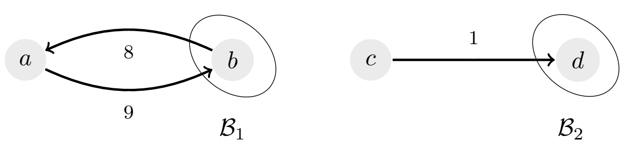

Figure 4: A visual explanation of Remark 5.1.

$\gamma \in S$

$\mathbb{D}_w$ *or* $\mathbb{D}_\top$ (Definition 14) $\longmapsto$ $\mathbb{D}_0$ *or* $\mathbb{D}_\perp$ [1]

$\mathbb{D}_{\bar{\gamma}}$ in Proposition 6

Figure 5: The range of γ-defence worsening γ from $\top$ to $\perp$, given a semiring $\mathbb{S} = \langle S, \oplus, \otimes, \perp, \top \rangle$.

In the following two propositions we respectively relate $\mathbb{D}_\gamma$ and $\mathbb{D}_2$ when all the arguments attack at most one other argument (Proposition 5), and when there are more attacks (Proposition 6).

Proposition 5. *Given $a \in \mathcal{A}_{rgs}$, we define T_a as $\bigcup_{b \in \mathcal{A}_{rgs}} R(a,b)$. If $\forall T_a$ the cardinality is $|T_a| \leqslant 1$, then $\mathbb{D}_2 \Leftrightarrow \mathbb{D}_\top$ (by Proposition 3, also $\mathbb{D}_2 \Leftrightarrow \mathbb{D}_w$ holds).*

Proof. If all $a \in \mathcal{A}_{rgs}$ attack one argument at most (i.e., $|T_a| \leqslant 1$), then $W(a, \mathcal{B} \cup \{b\})$ in Definition 14 is always equal to $W(a,b)$ in Definition 16. □

Given $\langle \mathcal{A}_{rgs}, R \rangle$, next proposition finds a $\bar{\gamma}$ such that $\mathbb{D}_2$ always implies $\mathbb{D}_{\bar{\gamma}}$.

Proposition 6 ($\mathbb{D}_2 \Rightarrow \mathbb{D}_{\bar{\gamma}}$)**.** *With T_a defined as in Proposition 5, if $\exists T_a . |T_a| \geqslant 2$, we find the n subsets T_a^i of T_a with cardinality $|T_a| - 1$. Then we define $\gamma_a = \bigsqcap_{i=1..n} (\bigotimes_{R(a,b) \in T_a^i} W(a,b))$, and $\bar{\gamma} = \bigsqcap \gamma_a$ ($\sqcap$ is the* glb *of $\mathbb{S}$). Finally, we obtain that $\mathbb{D}_2 \Rightarrow \mathbb{D}_{\bar{\gamma}}$ holds.*

Proof. The construction of γ_a finds, for every argument a, the worst composition $\bigotimes$ of all its subsets of attack weights with cardinality $|T_a| - 1$, where $|T_a|$ is the number of attacks from a. Hence, $W(a,b) \geqslant_{\mathbb{S}} W(\mathcal{B}, a) \Rightarrow W(a, \mathcal{B} \cup \{b\}) \oslash W(\mathcal{B}, a) \geqslant_{\mathbb{S}} \gamma_a$. If $\bar{\gamma} = \bigsqcap \gamma_a$, the previous implication is true for any $a \in \mathcal{A}_{rgs}$, and consequently $\mathbb{D}_2 \Rightarrow \mathbb{D}_{\bar{\gamma}}$. □

Finally, we can define an implication relation with respect to different γ:

Proposition 7. *If $\mathcal{B}$ γ_1-defends b and $\gamma_1 \geqslant_{\mathbb{S}} \gamma_2$, then $\mathcal{B}$ γ_2-defends b.*

Proof. If $(W(a, \mathcal{B} \cup \{b\}) \oslash W(\mathcal{B}, a)) \geqslant_{\mathbb{S}} \gamma_1$ and $\gamma_1 \geqslant_{\mathbb{S}} \gamma_2$, then $(W(a, \mathcal{B} \cup \{b\}) \oslash W(\mathcal{B}, a)) \geqslant_{\mathbb{S}} \gamma_2$. □

A graphical representation of γ-defence summarising Proposition 3, Proposition 4, and Proposition 6 and their ordering (Proposition 7) is given in Figure 5: by worsening γ we can switch from w-defence to Dung's defence [1] (at the two ends of Figure 5), and we can also exactly model $\mathbb{D}_2$ [2] for a given $\perp \leqslant_{\mathbb{S}} \bar{\gamma} \leqslant_{\mathbb{S}} \top$.

6. α^γ-semantics

In light of what advanced in the previous section, we are now ready to redefine some of the classical semantics in Abstract Argumentation by exploiting both the notion of *i)* an inconsistency amount α inside an extension (to be tolerated), and *ii)* the concept of γ-defence. In the following, for the sake of simplicity we will generically call these new semantics as α^γ-semantics. However, γ concerns all but the α-conflict-free semantics, since it is the only semantics (covered in this paper) not making use of the defence notion (directly or indirectly).

In Definition 18 we redefine the notion of conflict-freedom: conflicts can be now part of the solution up to a cost-threshold α. Such sets are now called α-conflict-free:

Definition 18 (α-conflict-free sets)**.** *Given a $WAAF_{\mathbb{S}}$, $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, a subset of arguments $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ is α-conflict-free iff $W(\mathcal{B}, \mathcal{B}) \geqslant_{\mathbb{S}} \alpha$.*

With respect to the $WAAF_{\mathbb{S}}$ in Figure 2, while the set $\{a, b, c\}$ is not conflict-free in the crisp version of the problem (since it includes the attacks between a and b, and between c and b), $\{a, b, c\}$ is instead 15-conflict-free because $W(a, b) + W(c, b) = 15$ (as a a reminder, we are using the weighted semiring).

Hence, by raising α we further relax the requirements behind conflict-freedom. This is highlighted by Proposition 8, since $\prod S = \bot$.

Proposition 8. *Given any $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, the set of $\bot$-conflict free sets correspond to the power-set of $\mathcal{A}_{rgs}$.*

No constraint is given on the amount of conflict internal to a set, thus all the arguments can coexist together.

We now define two propositions that derive from Definition 18 and from the semiring properties explained in Section 2.2.

Proposition 9. *If a set is α_1-conflict-free and $\alpha_1 \geqslant_{\mathbb{S}} \alpha_2$, then the same set is also α_2-conflict-free.*

Proof. If $W(\mathcal{B}, \mathcal{B}) \geqslant_{\mathbb{S}} \alpha_1$ and $\alpha_1 \geqslant_{\mathbb{S}} \alpha_2$, then $W(\mathcal{B}, \mathcal{B}) \geqslant_{\mathbb{S}} \alpha_2$. □

For instance, $\{a, b, c\}$ is also 16-conflict-free because it is a 15-conflict-free ($15 \geqslant_{\mathbb{S}} 16$ in the weighted semiring).[5] Therefore, this states than in α-conflict-free sets we tolerate an internal inconsistency-amount better than α.

We now introduce α^γ-admissible sets:

Definition 19 (α^γ-admissible sets)**.** *Given $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, an α-conflict-free set $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ is α^γ-admissible iff the arguments in $\mathcal{B}$ are γ-defended by $\mathcal{B}$ from the arguments in $\mathcal{A}_{rgs} \backslash \mathcal{B}$.*

Considering the framework in Figure 2 as unweighted (i.e., the one in Figure 1), Dung's admissible sets are: $\{a\}, \{c\}, \{d\}, \{a, c\}, \{a, d\}$. $\top^\top$-admissible sets (i.e., 0^0-sets in the weighted semiring) are $\{a\}$, $\{c\}$, and $\{a, c\}$ instead: $\{a\}$ because is not attacked by any other argument, $\{c\}$ and $\{a, c\}$ because they both w-defends c from the attack performed by d, i.e., $W(d, c) \geqslant_{\mathbb{S}} W(c, d)$ (i.e., $8 \leqslant 9$). For instance $\{d\}$ is not 0^0-admissible because it is not able to 0-defend

[5] In the weighted semiring, $\leqslant_{\mathbb{S}}$ is equivalent to $\geqslant$ over Real numbers, while in the probabilistic and fuzzy semirings, $\geqslant_{\mathbb{S}}$ corresponds to $\geqslant$ over Real numbers in the interval [0..1].

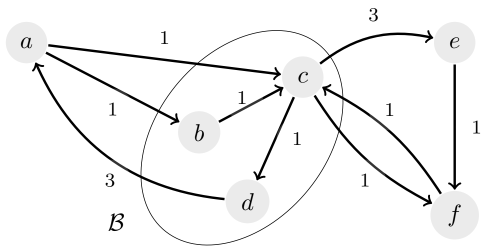

Figure 6: An example of 2^0-complete extension, $\mathcal{B} = \{b, c, d\}$; $\mathcal{B} \cup \{f\}$ is 4^0-complete, while $\mathcal{B} \cup \{f\}$ and $\mathcal{B} \cup \{e\}$ are two 5^0-complete extensions (using the weighted semiring).

(or w-defend, see Proposition 3) itself from the attack of c. For the same reason, $\{a, d\}$ is not 0^0-admissible. As we can see from this example, w-defence restrain Dung's defence, and, accordingly the number of sets satisfying admissibility: in this case we drop $\{d\}$ and $\{a, d\}$ w.r.t. [1].

Considering an example with an internal inconsistency $\alpha \neq \top$, the set $\{a, b, c\}$ is 15^0-admissible: it is 15-conflict-free, and $\{a, b, c\}$ 0-defends its arguments, i.e., c from d. All the 15^0-admissible sets are $\emptyset$, $\{c\}$, $\{c, e\}$, $\{a\}$, $\{a, c\}$, $\{a, c, e\}$, and $\{a, b, c\}$. In order to provide an example with both $\alpha \neq \top$ and $\gamma \neq \top$ (still considering Figure 2), the set $\{d, e\}$ is 11^1-admissible, since it is 11-conflict-free, and d defends itself (and the whole $\{d, e\}$) from c by paying a penalty of $9 - 8 \leqslant 1$.

Three semantics that refine α^γ-admissibility are introduced from Definition 20 to Definition 22:

Definition 20 (α^γ-complete semantics)**.** *Given $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S} \rangle$, an α^γ-admissible $\mathcal{B} \subseteq \mathcal{A}_{rgs}$ is α^γ-complete iff each argument $b \in \mathcal{A}_{rgs}$ that is γ-defended by $\mathcal{B}$ and s.t. $W(\mathcal{B} \cup \{b\}, \mathcal{B} \cup \{b\}) \geqslant_{\mathbb{S}} \alpha$ is in $\mathcal{B}$ (i.e., $b \in \mathcal{B}$).*

For instance, the set $\{a, d\}$ in Figure 2 is complete according to [1] (i.e., not considering weights), but it is not 0^0-complete because d is not able to w-defend $\{a, d\}$ from c. However, $\{a, d\}$ is 0^1-complete by considering the weighted semiring and allowing $\gamma = 1$. Note that $\{a, d\}$ also 1-defends argument e, but $\{a, d, e\}$ is not 0^1-complete if we keep $\alpha = 0$: bringing e inside would lead to an internal conflict of $W(d, e) + W(e, e) = 11$.

Therefore, in the α^γ-complete semantics we need to bring in all the γ-defended arguments while respecting the α-threshold at the same time. An example is given in Figure 6 (we suppose to use the weighted semiring), where $\mathcal{B} = \{b, c, d\}$ is the only 2^0-complete extension: even if $\mathcal{B}$ 0-defends f from e, it is not possible to bring f in $\mathcal{B}$ because we can tolerate only 2 as internal conflict (and $W(b, c) + W(c, d) = 2$ already). However, by relaxing the problem to find 4^0-complete extensions, $\{b, c, d, e\}$ is the single solution, while both $\{b, c, d, e\}$ and $\{b, c, d, f\}$ are two 5^0-complete extensions.

Definition 21 (α^γ-preferred)**.** *An α^γ-preferred extension is a maximal (with respect to set inclusion) α^γ-admissible subset of $\mathcal{A}_{rgs}$.*

Still considering Figure 2, $\{a, c\}$ and $\{a, d\}$ are the two preferred extensions according to [1] (i.e., not considering weights). However, $\{a, c\}$ is the only 0^0-preferred extension, while $\{\{a, c\}, \{a, d\}\}$ is the set of 0^1-preferred extensions.

Definition 22 proposes Dung's stable semantics revisited in a $WAAF_{\mathbb{S}}$.

Definition 22. **(α^γ-stable semantics)** *Given $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, an α^γ-admissible set $\mathcal{B}$ is also an α^γ-stable extension iff $\forall a \notin \mathcal{B}, \exists b \in \mathcal{B}. W(b, a) \neq \top$, and $\mathcal{B} \cup \{a\}$ is not α^γ-admissible.*

Therefore, each argument a left outside $\mathcal{B}$ needs to be attacked ($W(b, a) \neq \top$), and it is not possible to bring a in $\mathcal{B}$ without exceeding α as internal inconsistency. For example, considering the problem in Figure 2 as unweighted (i.e., as a classical AAF), the set $\{a, d\}$ corresponds to the only stable extension. However, considering weights, this set is not 0^0-stable, because $W(c, d) \not\geqslant_{\mathbb{S}} W(\{a, d\}, c)$, i.e., $9 \nleqslant 8$. However, it is 0^1-stable, since $W(c, d) \oslash W(\{a, d\}, c) = 9 - 8 \leqslant 1$ satisfies $\gamma = 1$. Thus, in such example there is no 0^0-stable extension.

7. Some Formal Results on α^γ-semantics

In this section we provide general considerations on α-semantics, as their properties and inclusion relations. The first result we present is that classical inclusion-relations [1] among the newly-defined α-semantics are still valid:

Theorem 7.1 (α^γ-semantics inclusions)**.** *Given any $\langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, with $\mathbb{S} = \langle S, \oplus, \otimes, \bot, \top\rangle$, and $\alpha, \gamma \in S$,*

1. *each α^γ-admissible set is also α-conflict-free.*
2. *each α^γ-complete extension is also an α^γ-admissible set.*
3. *each α^γ-preferred extension is also α^γ-complete.*
4. *each α^γ-stable extension is also α^γ-preferred.*

Proof. *1*) and *2)* are straightforwardly proved by definition (see Definition 18, Definition 19, and Definition 20): to be α-conflict-free (or, α^γ-admissible) is a necessary condition to be also α^γ-admissible (or, α^γ-complete). *3)*, if $\mathcal{B}$ is a maximal set such that each argument in $\mathcal{B}$ is γ-defended by $\mathcal{B}$, then each argument which is γ-defended by $\mathcal{B}$ is in $\mathcal{B}$ (i.e., $\mathcal{B}$ is also α^γ-complete). Concerning *4)*, by definition an α^γ-stable extension is also an α^γ-admissible set (see Definition 22); in addition, it is also maximal with respect to set inclusion because it is not possible to bring one more element inside (without exceeding the internal inconsistency-threshold given by α). □

Theorem 7.1 leads to Corollary 7.2, which states that the classical implication chain between semantics [1] also holds for α-semantics.

Corollary 7.2. *By setting $\alpha, \gamma \in S$, the following implications hold between α^γ-semantics and sets: α^γ-stable $\Rightarrow$ α^γ-preferred $\Rightarrow$ α^γ-complete $\Rightarrow$ α^γ-admissible $\Rightarrow$ α-conflict-free.*

Theorem 7.3 shows when α^γ-semantics can be used to exactly obtain the original semantics [1].

Theorem 7.3. *Given $F = \langle \mathcal{A}_{rgs}, R\rangle$, and $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, with $\mathbb{S}$ as desired, then*

1. *the set of $\top$-conflict-free sets in WF is equal to the set of conflict-free sets in F.*
2. *the set of $\top^\perp$-admissible sets in WF is equal to the set of admissible sets in F.*
3. *the set of $\top^\perp$-complete extensions in WF is equal to the set of complete extensions in F.*
4. *the set of $\top^\perp$-preferred extensions in WF is equal to the set of preferred extensions in F.*
5. *the set of $\top^\perp$-stable extensions in WF is equal to the set of stable extensions in F.*

Proof. The proof can be straightforwardly achieved by considering that $\alpha = \top$ leads to not allowing any attack in an extension (as in [1]), and $\perp$-defence is equivalent to the original definition of defence [1] (see Proposition 4). □

Theorem 7.4 relates α^γ-semantics using a non-relaxed defence (i.e., w-defence or $\top$-defence) and no internal conflict (i.e., $\alpha = \top$) to their counterpart in the classical ones [1]. The result is the semantics we have presented in [8]. Note that the results obtained in Theorem 7.4 differ from those in Theorem 7.3. Having $\gamma = \top$ noticeably impacts on the hierarchy obtained in Theorem 7.3, since four out of five points change between the two theorems.

Theorem 7.4. *Given $F = \langle \mathcal{A}_{rgs}, R\rangle$, and $WF = \langle \mathcal{A}_{rgs}, R, W, \mathbb{S}\rangle$, with $\mathbb{S}$ as desired, then*

1. *the set of $\top$-conflict-free sets in WF is equal to the set of conflict-free sets in F.*
2. *the set of $\top^\top$-admissible sets in WF is a subset of the set of admissible sets in F.*
3. *for each $\top^\top$-complete extension $\mathcal{B}_{WF}$ in WF, there exists a complete extension $\mathcal{B}_F$ in F, s.t., $\mathcal{B}_{WF} \subseteq \mathcal{B}_F$.*
4. *for each $\top^\top$-preferred extension $\mathcal{B}_{WF}$ in WF, there exists a preferred extension $\mathcal{B}_F$ in F, s.t. $\mathcal{B}_{WF} \subseteq \mathcal{B}_F$.*
5. *for each $\top^\top$-stable extension $\mathcal{B}_{WF}$ in WF, there exists a stable extension $\mathcal{B}_F$ in F, s.t. $\mathcal{B}_{WF} \subseteq \mathcal{B}_F$.*

Proof. The two notions of w-defence and γ-defence are interchangeable when $\gamma = \top$ (see Proposition 3). Concerning *1)*, a semiring-value equal to $\top$ represents a no-attack relation, so $\top$-conflict-free sets do not include any attack. *2)* holds because the notion of w-defence implies the classical notion of defence ($\mathbb{D}_w \Rightarrow \mathbb{D}_0$), but not vice versa (see Proposition 1). *4)* follows from Definition 6, and Definition 21: the maximal $\top^\top$-admissible sets w.r.t. set inclusion are computed over a subset of the admissible ones ($\mathbb{D}_w \Rightarrow \mathbb{D}_0$ in Proposition 1). Therefore, each of them is a subset of at least one preferred extension in the corresponding unweighted framework. In order to visually understand the proof of this item, Figure 7 shows an example of inclusion hierarchy among admissible sets, where maximal ones (i.e., preferred extensions) are G, H, E, I, L. Since $\top^\top$-admissible sets are a subset of admissible sets (see Theorem 7.4), the inclusion hierarchy in Figure 7 can change to Figure 8. Accordingly, each $\top^\top$-preferred extension (i.e., B, I) is a subset of a preferred extension. Same considerations hold for *3)* and *5)*, since $\mathbb{D}_w \Rightarrow \mathbb{D}_0$ less arguments need to be taken in order to have a valid $\top^\top$-complete extension. □

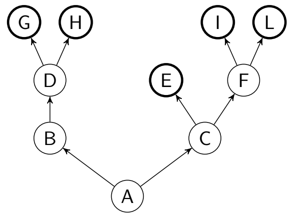

Figure 7: An example of inclusion hierarchy among admissible sets: maximal, i.e., preferred ones, are G, H, E, I, L.

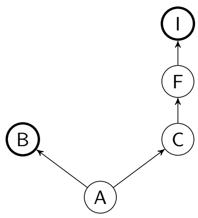

Figure 8: Since $\top^\top$-admissible sets are a subset of admissible ones, the hierarchy in Figure 7 becomes as in this figure. $\top^\top$-preferred extensions are B and I.

Theorem 7.5 shows what happens to α-semantics when α and γ are worsened.

Theorem 7.5. *Given $\langle \mathcal{A}_{rgs}, R, W, S = \langle A, \oplus, \otimes, \bot, \top \rangle \rangle$, and $\alpha_1, \alpha_2, \gamma_1, \gamma_2 \in A$ s.t. $\alpha_1 \geqslant_{\mathbb{S}} \alpha_2$ and $\gamma_1 \geqslant_{\mathbb{S}} \gamma_2$, then*

1. *the set of α_1-conflict-free sets is a subset of the set of α_2-conflict-free sets,*
2. *the set of $\alpha_1^{\gamma_1}$-admissible sets is a subset of the set of $\alpha_2^{\gamma_2}$-admissible sets,*
3. *for each $\alpha_1^{\gamma_1}$-complete extension $\mathcal{B}_1$, there exists an $\alpha_2^{\gamma_2}$-complete extension $\mathcal{B}_2$, such that $\mathcal{B}_1 \subseteq \mathcal{B}_2$.*
4. *for each $\alpha_1^{\gamma_1}$-preferred extension $\mathcal{B}_1$, there exists an $\alpha_2^{\gamma_2}$-preferred extension $\mathcal{B}_2$, such that $\mathcal{B}_1 \subseteq \mathcal{B}_2$.*
5. *for each $\alpha_1^{\gamma_1}$-stable extension $\mathcal{B}_1$, there exists an $\alpha_2^{\gamma_2}$-stable extension $\mathcal{B}_2$, such that $\mathcal{B}_1 \subseteq \mathcal{B}_2$.*

Proof. Points *1)* and *2)* can be proved by using Proposition 9 and Proposition 7. To prove *3)*, *4)*, and *5)* is identical to prove the respective items in Theorem 7.3, since $\alpha_2^{\gamma_2}$-semantics represent a relaxation of $\alpha_1^{\gamma_1}$-semantics, as classical semantics in [1] are a relaxation of $\top^\top$-semantics. □

From Theorem 7.5 we can also derive some results on the justification of arguments (see Definition 8) when α and γ are increased.

Corollary 7.6. *Given* $\langle \mathcal{A}_{rgs}, R, W, S = \langle A, \oplus, \otimes, \bot, \top \rangle \rangle$, $\alpha_1, \alpha_2, \gamma_1, \gamma_2 \in A$ *s.t.* $\alpha_1 \geqslant_{\mathbb{S}} \alpha_2$ *and* $\gamma_1 \geqslant_{\mathbb{S}} \gamma_2$, *and given any semantics* $\sigma \in \{adm, com, prf, stb\}$, *the set of defensible arguments in* $\sigma_{\alpha_1}^{\gamma_1}$ *is a subset of the defensible arguments in* $\sigma_{\alpha_2}^{\gamma_2}$.

Proof. The proof consists in showing that the union of all the extensions with a given semantics σ is non-decreasing if α and/or γ increase: $\bigcup \mathcal{E}_{\sigma_{\alpha_1}^{\gamma_1}} \subseteq \bigcup \mathcal{E}_{\sigma_{\alpha_2}^{\gamma_2}}$. This is trivially true for item 2 in Theorem 7.5. In addition, it holds also for item 3, item 4, and item 5: for all $\mathcal{B}_1$ there exists $\mathcal{B}_2$ such that $\mathcal{B}_1 \subseteq \mathcal{B}_2$, hence this inclusion also holds for their union. □

8. Implementation and Tests

We have implemented α^γ-semantics in *ConArg*[6] [14], which is a tool that exploits *Gecode*[7] (a constraint-programming library) to solve several problems related to Argumentation. All the following tests have been collected on a benchmark of 100 graphs (25 arguments each) generated according to the Erdős-Rényi random model [15]: a generator in the *NetworkX* library[8] has been used. Each directed edge is added to a graph with an independent probability p. With each edge we associate a random natural number in the interval [1..10] (in order to test $\mathbb{S}_{weighted}$), and [1..10]/10 (to test $\mathbb{S}_{fuzzy}$), using a uniform distribution.

Figure 9 and Figure 10 respectively show the average number (on 100 graphs) of α^γ-admissible and α^γ-stable extensions (other semantics are omitted for the sake of space) for all the 78 combinations of $\alpha = \{0, 1, 2, 4, 6, 8, 9, 10, 11, 12\}$ and $\gamma = \{0, 1, 2, 4, 6, 8\}$, using $\mathbb{S}_{weighted}$: from these figures, we can see what happens when both α and γ change. The two sets]grow in the same way, even if they reach a cardinality of 525 and 21 respectively.

Figure 11 shows the same using $\mathbb{S}_{fuzzy}$, for all the 121 combinations of $\alpha, \gamma = \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$. Figure 12 presents the same results in Figure 11 with a focus on the interval $\alpha, \gamma = \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$ to show what happens before reaching the maximum relaxation for both α, γ (i.e., these images do not show values when $\alpha, \gamma \in \{0, 0.1\}$)

Figure 13, Figure 14, Figure 15, and Figure 16 visually present the result in Corollary 7.6: the set of defensible arguments is non-decreasing while increasing α and/or γ. Figure 13 (α^γ-admissible) and Figure 14 (α^γ-stable) show this for

[6] `http://www.dmi.unipg.it/conarg/`.
[7] `http://www.gecode.org`.
[8] `https://networkx.github.io`.

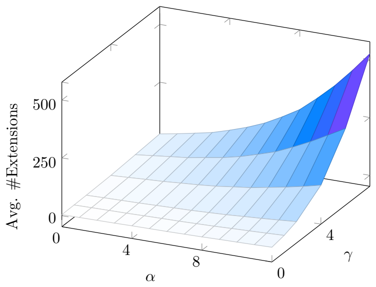

Figure 9: Average number of α^γ-admissible sets (changing α and γ); with $\mathbb{S}_{weighted}$.

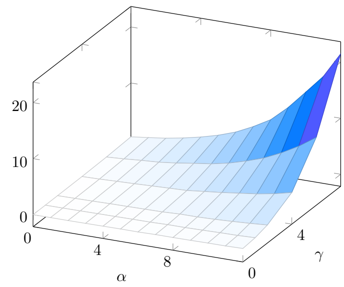

Figure 10: Average number of α^γ-stable extesions (changing α and γ); with $\mathbb{S}_{weighted}$.

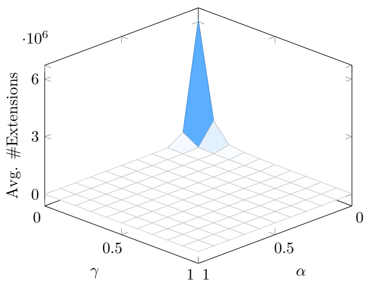

Figure 11: Average number of α^γ-admissible sets (changing α and γ); with $\mathbb{S}_{fuzzy}$.

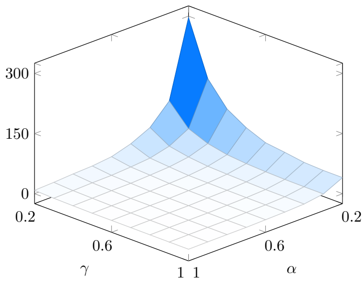

Figure 12: Same results in Figure 11 with a focus on the interval $\alpha, \gamma = \{0.2, \ldots, 1\}$.

$\mathbb{S}_{weighted}$, while Figure 15 (α^γ-admissible) and Figure 16 (α^γ-stable) show the same for $\mathbb{S}_{fuzzy}$.

Figure 17 reports instead the average number of α^γ-admissible sets, α^γ-complete, α^γ-preferred, and α^γ-stable extensions using $\mathbb{S}_{weighted}$ (Figure 17a and Figure 17b) and $\mathbb{S}_{fuzzy}$ (Figure 17c and Figure 17d). In Figure 17a and Figure 17c we change α while we keep $\gamma = \top$. In Figure 17b and Figure 17d we change γ while we keep $\alpha = \top$. From these figures we see that semantics implications are respected (see also Corollary 7.2). We also see that the number of extensions remains quite stable (except for α^γ-admissible sets). Note that a report of (time) performance is outside the scope of this work: however, for instance, the average time to find all $\top^\perp$-admissible sets is very close to the time to find all admissible sets.

9. Applications

In this section we present two different applications of our framework. The first one (Section 9.1) concerns removing the effect of troll arguments: the exam-

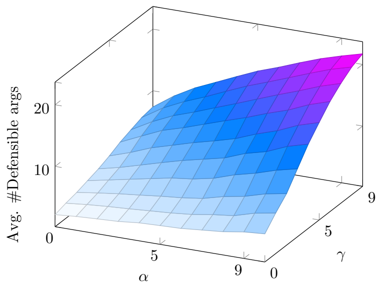

Figure 13: Average defensible arguments in α^γ-admissible sets ($\mathbb{S}_{weighted}$).

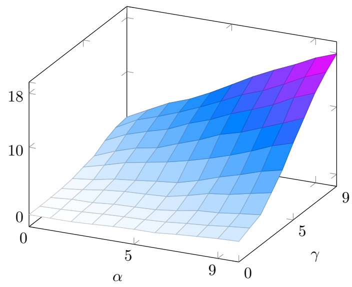

Figure 14: Average defensible arguments in α^γ-stable extensions ($\mathbb{S}_{weighted}$).

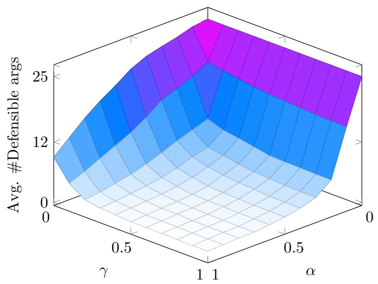

Figure 15: Average defensible arguments in α^γ-admissible sets ($\mathbb{S}_{fuzzy}$).

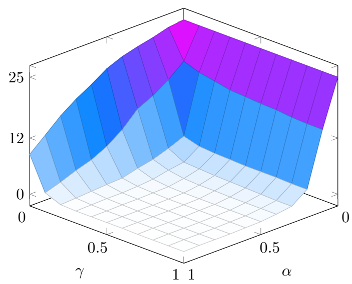

Figure 16: Average defensible arguments in α^γ-stable extensions ($\mathbb{S}_{fuzzy}$).

ple is directly taken from the work in [5]. The second application (Section 9.2) concerns a real-world scenario based on arguments taken from an Amazon.com review. A different approach to this case-study has been already presented in [16], where the goal was to show that negative arguments increasingly permeate positive reviews over time.

9.1. Trolls

Relaxing a framework allows us to mitigate the disturbing effect of poorly specified or unsound attacks (*e.g.*, from *trolls*) [5]. In Figure 18 we show exactly the same framework (with the same weights) reported in [5], where several participants argue about the role of the government in what banning smoking is concerned. Arguments are:

a. *Governments should ban smoking.*

b. *Governments shan't interfere with the right to smoke.*

c. *Smoking is a matter of freedom of choice and governments ought to protect the rights of their citizens. Therefore, banning smoking would be a violation of rights.*

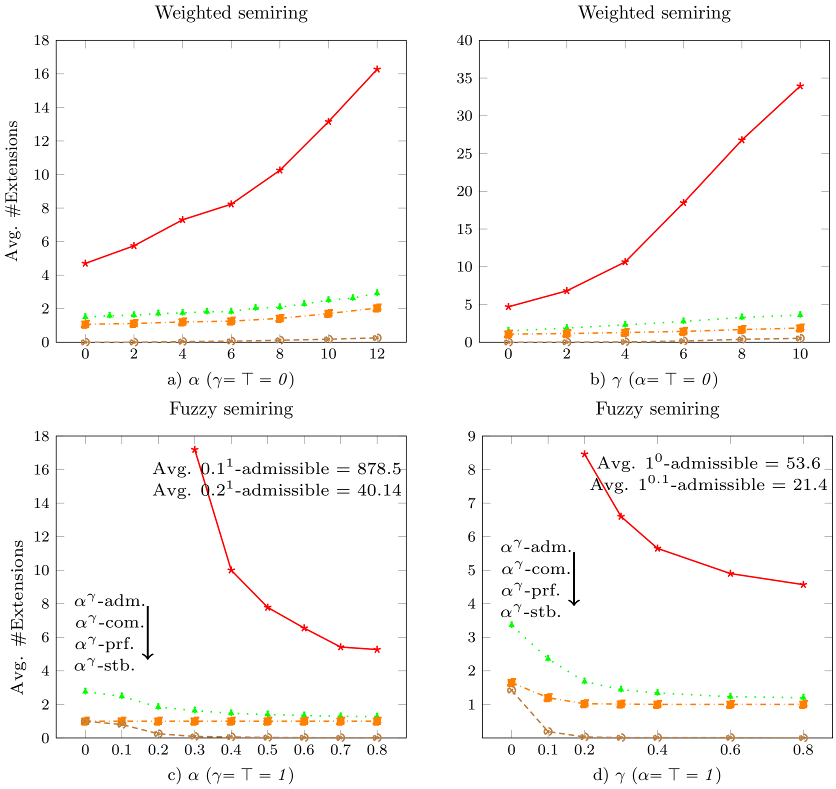

Figure 17: α^γ-admissible ( ), -complete ( ), -preferred ( ), and -stable ( ) number of average sets/extensions (on 100 graphs), using the Fuzzy semiring.

d. *Time after time, clinical research has proved that smoking is highly addictive. Thus, the issue may not be considered as a matter of freedom of choice, and governments are supposed to ban these practices.*

e. *I like turtles.*

Weights in Figure 18 represent a strength score for each attack. The attack from e to a is meant to represent a troll attack (its strength is very low, i.e., 1). In [5] the authors show how their computational framework is able to mitigate the disturbing effect of such attack: e is not attacked, thus, with the classical semantics [1], it is capable of always ruling a out (its impact is strong).

However, we can mitigate it also by using our framework: if we compute the $\top^\perp$-stable extensions (by using $\gamma = \perp$ we consider classical defence [1]) we obtain $\{d, e\}$ as the sole solution: in [5] these are the same two most preferred arguments before mitigating the troll attack. If we instead relax the problem

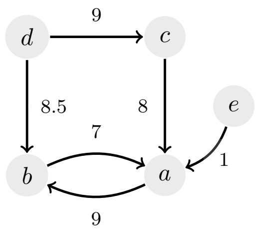

Figure 18: The *troll* framework example in [5], using $\mathbb{S}_{weighted}$.

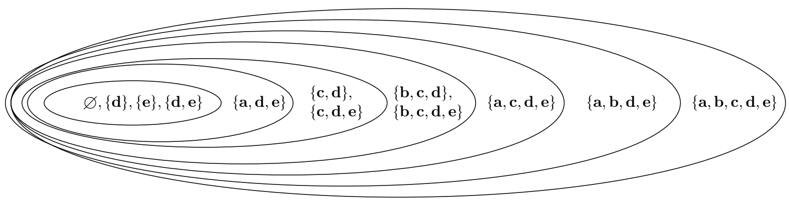

Figure 19: The α to obtain the represented $\alpha^\top$-admissible sets on the WAAF in Figure 18 is $0, 1, 9, 17, 18, 25, 42$, from inner to outer sets.

by computing the $1^\perp$-stable semantics, the solution becomes $\{a, d, e\}$. This extension contains the same three most preferred arguments in [5] (i.e., a, d, and e) after mitigating the troll attack: thus, we remove the effect of e on a in the framework. Note that e is a "good" argument in [5], thus it is not surprising it can be part of a "good" extension. It is the attack from e to a to be fake: the aim is to remove the effect of the attack from e, not e itself.

In Figure 19 we show as the slices of an onion the set of all $\alpha^\perp$-admissible sets that can be obtained over Figure 18: $\{\emptyset, \{d\}, \{e\}, \{d,\ e\}, \{a, d, e\}, \{c, d\}, \{c, d, e\}, \{b, c, d\}, \{b, c, d, e\}, \{a, c, d, e\}, \{a, b, d, e\}, \{a, b, c, d, e\}\}$, which can be obtained by varying $\alpha = 0, 1, 9, 17, 18, 25, 42$ respectively. The inner slice contains only $0^\top$-admissible sets; all the other sets respectively contain (from inner to outer) $1^\top$-admissible, $9^\top$-admissible, $17^\top$-admissible, $18^\top$-admissible, $25^\top$-admissible, and $42^\top$-admissible sets. $0^\top$-admissible are also $1^\top$-admissible sets, and so on. By fixing $\alpha = \top$ and changing γ for the same example in Figure 19, it is instead not possible to obtain more slices as for changing only α, since the admissible sets are always $\{\emptyset, \{e\}, \{d\}, \{d, e\}\}$, as for $\alpha = \top$.

In Figure 21 we show instead all the sets of $\top^\gamma$-admissible sets with respect to the WAAF in Figure 20, that is we only vary γ by keeping $\alpha = \top$ (i.e., not allowing internal inconsistency). We obtain the same onion-slices representation in Figure 19 (where instead $\gamma = \top$ and α changes). If $\gamma = 0$ (i.e., using w-defence), then the only $\top^\top$-admissible sets are $\emptyset$ and $\{a\}$. If $\gamma = 1$, $\{b\}$ is the single $\top^1$-admissible set. If $\gamma = 2$, the $\top^2$-admissible set is $\{c\}$, while if $\gamma = 3$ the $\top^3$-admissible sets are $\{d\}$ and $\{b, c\}$. Then, $\{b, d\}$ is $\top^4$-admissible, $\{c, d\}$ is $\top^5$-admissible, and, finally, $\{b, c, d\}$ is $\top^6$-admissible. By only further increasing

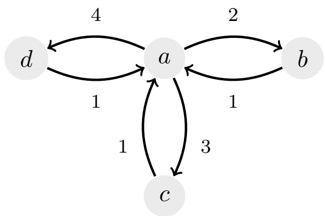

Figure 20: An example to show how changing γ (with $\alpha = \top$), $\top^\gamma$-admissible sets change (see Figure 21).

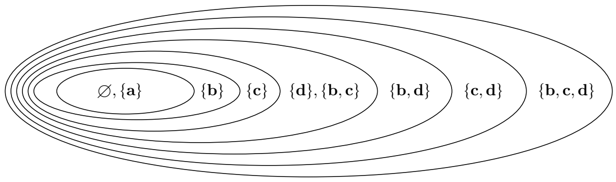

Figure 21: The γ to obtain the represented $\top^\gamma$-admissible sets on the WAAF in Figure 20 is $0, 1, 2, 3, 4, 5, 6$, from inner to outer sets.

γ it is not possible to obtain more admissible sets.

Finally, Figure 22 is presented to show how internal and defence relaxations are strictly linked together: the set $\{a, d\}$ is $\top^3$-admissible, since a is attacked by c with weight of 8, but only a counter-attack with weight 5 is present from d to c; hence, in the weighted semiring, the difference to be tolerated is $8 - 5 = 3$. However, if an internal inconsistency of 2 can be tolerated, the set $\{a, d, e\}$ is $2^\top$-admissible: by allowing a small internal conflict, the defence of $\{a, d, e\}$ becomes stronger, since no γ relaxation is needed. Therefore, we provide a means to an agent to decide between $\{a, d\}$, which has no internal conflict, or $\{a, d, e\}$, which has a stronger defence: this naturally leads to defining a two-criteria ranking, whose investigation is however outside the scope of this work (see future work in Section 11).

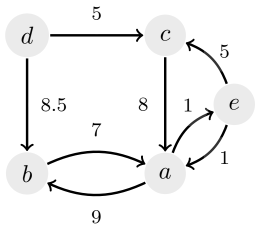

Figure 22: $\{a, d\}$ is $\top^3$-admissible, $\{a, d, e\}$ is $2^\top$-admissible.

9.2. A Study on an Amazon.com Review

The second application is based on a set of arguments extracted from an *Amazon.com* review. Since automated tools performing *argument mining* [17] still need to be refined to obtain excellent results, and since we need to extract the whole WAAF also including attacks, we manually extracted all the arguments from each review on one chosen product, and the attacks among them.

For our purposes, we retrieved the Amazon "Clothing and Accessories" section from the SNAP project database.[9] This dataset contains approximately 110k products and spans from 1999 to July 2014, for a total of more than one million reviews. We looked for a product that had two characteristics: *(a) it has a quite large number of reviews*; *(b)* such number of reviews is not too large to be tracked down by hand. We randomly extracted products from the list, discarding those that did not fit the bill. Eventually, we came across one that fitted and which happened to be a ballet tutu for kids. We examined all of the 240 reviews that this product received between 2009 and July 2014. From the reviews, we collect a total of 24 positive (in favour of purchase) arguments and 20 negative (against purchase) arguments, whose absolute frequencies are reported in Table 1. Note we also aggregated arguments with the same meaning into a single argument: at the moment, we are not aware of any automated tool performing this aggregation step. The number of repetitions represents a strength valued for that argument.

For some couples of arguments, posing attacks has been very easy: some positive arguments are the exact negation of what stated in the relative negative argument (and vice versa). For instance, looking at Table 1, *the tutu has a good quality* (D) and *the tutu has a bad quality* (a), or *the tutu fits well* (B) and *the tutu does not fit* (c). For the sake of completeness, such easy-to-detect (bidirectional) attacks are $\{B \leftrightarrow c, C \leftrightarrow l, D \leftrightarrow a, E \leftrightarrow f, F \leftrightarrow h, I \leftrightarrow o, J \leftrightarrow d, M \leftrightarrow e, N \leftrightarrow k, P \leftrightarrow m, S \leftrightarrow b, T \leftrightarrow g, X \leftrightarrow j\}$. Note that we also have two unidirectional attacks between two positive arguments ($Q \rightarrow N$ and $V \rightarrow J$), and one bidirectional attack between two negative arguments ($s \leftrightarrow h$). Other unidirectional and bidirectional attacks are $p \rightarrow F$, $c \rightarrow O$, $c \rightarrow G$, $d \rightarrow O$, $d \rightarrow G$, $M \rightarrow d$, $M \rightarrow o$, $e \rightarrow O$, $e \rightarrow J$, $e \rightarrow I$, $t \rightarrow P$, $j \rightarrow O$, $i \rightarrow D$, $S \rightarrow r$, $S \rightarrow a$, $a \rightarrow N$, $a \rightarrow O$, $U \rightarrow b$, $n \rightarrow U$, $E \rightarrow a$, $o \rightarrow O$, $o \rightarrow G$, $s \leftrightarrow F$, $f \leftrightarrow L$, $f \leftrightarrow K$, for a total of 58 attacks. The weight of each attack is represented by the number of repetitions of its tail-argument in Table 1. The full WAAF, on which we work by using the weighted semiring, can be found in [16].

On this WAAF we obtain 256 stable extensions, or equivalently 0^{∞}-stable extensions are 256 (see Theorem 7.3): this is quite a large number, due to the fact that most of the attacks are symmetric, and for this reason it is not easy to extract some results from such a large set of extensions.[10] For this reason, it can be useful to apply a more sceptical approach, which can be provided by

[9]Courtesy of Julian McAuley and SNAP project (source: `http://snap.stanford.edu/data/web-Amazon.html` and `https://snap.stanford.edu`).

[10]Clearly, preferred or other extensions are even more in number.

Table 1: Positive and negative arguments, with their number of appearances in reviews between 2009 and July 2014.

ID	Positive arguments	#App.	ID	Negative arguments	#App.
A	the kid loved it	78	a	it has a bad quality	18
B	it fits well	65	b	it is not sewed properly	17
C	it has a good quality/price ratio	52	c	it does not fit	12
D	it has a good quality	44	d	it is not full	11
E	it is durable	31	e	it is not as advertised	8
F	it is shipped fast	25	f	it is not durable	7
G	the kid looks adorable	23	g	it has a bad customer service	4
H	it has a good price	21	h	it is shipped slow	3
I	it has great colors	21	i	it smells chemically	3
J	it is full	18	j	you can see through it	3
K	it did its job	11	k	it cannot be used in real dance class	2
L	it is good for playing	11	l	it has a bad quality/price ratio	2
M	it is as advertised	9	m	it has a bad envelope	1
N	it can be used in real dance classes	7	n	it has a bad waistband	1
O	it is aesthetically appealing	7	o	it has bad colours	1
P	it has a good envelope	2	p	it has high shipping rates	1
Q	it is a great first tutu	2	q	it has no cleaning instructions	1
R	it is easier than build your own	2	r	it is not lined	1
S	it is sewed properly	2	s	it never arrived	1
T	it has a good customer service	1	t	it was damaged	1
U	it is secure	1			
V	it is simple but elegant	1			
W	you can customize it	1			
X	you cannot see through it	1			

tuning α and γ in a different way. At first, we look for 0^0-stable extensions, trying to find stable extensions for which the defence weight of an extension is greater or equal than the attack weight from outside to inside; however, we find there is no subset of arguments with such a property on the considered WAAF. Thus, we turn our attention to find the stable extensions with the best γ (i.e., lowest) while keeping $\alpha = 0$ (i.e., no internal conflict as in [1]): in this way, we obtain 16 different 0^{22}-stable extensions.

The final step consists in also allowing a small internal conflict (increasing α) in order to reduce the relaxation on the defence (decreasing γ). We can find only a couple of 1^7-stable extensions, which correspond to:

$$\{A, B, C, E, F, G, H, I, K, L, M, Q, R, T, V, W, b, i, j, k, m, n, p, q, r, t\}$$

$$\{A, B, C, E, F, G, H, I, K, L, M, Q, R, V, W, b, g, i, j, k, m, n, p, q, r, t\}$$

This represents the most sceptical approach preserving some result: by further reducing either α or γ we do not obtain α^γ-stable extensions anymore. Therefore, by only allowing a small conflict of 1 (p is repeated only once in Table 1), which corresponds to the attack $p \rightarrow F$, we find two stable extensions that defend themselves better (i.e., with a lesser γ) than the others, with only $\gamma = 7$. Since only one reviewer out of 240 thinks shipping rates are expensive (argument p), then including the conflict $p \rightarrow F$, where fast delivery of the product (argument F) has been experienced by 25 reviewers instead, can be tolerable in order to obtain stable extensions with a stronger defence.

10. Related Work and Comparison

Besides the technical comparison between the most related definitions of weighted defence [2, 3] (see Section 4), in this section we introduce other frameworks in the literature, which take into account weights or preference-values on either arguments or attacks.

In [18], AAFs have been extended to *Value-based* AAFs (*VAFs*). A VAF is a five-tuple $\langle \mathcal{A}_{rgs}, R, V, val, P \rangle$, where $\mathcal{A}_{rgs}$ is a finite set of arguments, R is an irreflexive binary relation on A (i.e., $\langle \mathcal{A}_{rgs}, R \rangle$ is a standard AAF), V is a non-empty set of values, val is a function which maps from elements of A to elements of V, and P is the set of possible audiences (i.e., total orders on V). We say that an argument a relates to value v if accepting A promotes or defends v: the value in question is given by *val(a).* For every $a \in \mathcal{A}_{rgs}$, $val(a) \in V$. When the VAF is considered by a particular audience, the value ordering is fixed. A *Preference-based* argumentation AAF [19] is a triplet $\langle \mathcal{A}_{rgs}, R, Pref \rangle$ where *Pref* is a partial pre-ordering (reflexive and transitive binary relation) on $\mathcal{A}_{rgs} \times \mathcal{A}_{rgs}$. The notion of defence changes accordingly: let a and b be two arguments, b attacks a iff $R(b, a)$ and not $a > b$, i.e., a is not preferred in the partial pre-ordering.

In [20] the author extends Dung's theory of argumentation to integrate a meta-level argumentation concerning preferences. Dung's level of abstraction is preserved, so that arguments expressing preferences are distinguished by being the source of a second attack relation. This abstractly characterises the application of preferences by attacking attacks between the arguments that are subject to preference claims.

Hence, the first three presented references "weigh" arguments instead of attacks, and, moreover, the proposed frameworks are qualitatively-oriented instead of quantitatively-oriented (as in our proposal). To comment on the first issue, we remind that it can be possible to aggregate the weights on the attacks in order to obtain a problem with preferences/scores: examples of such approaches are explained in [7, 20, 21]. Note that the semiring operator $\oplus$ can represent both a partial [19] and a total [18] order among the arguments. In addition, the $\otimes$ operator of the semiring can be used to aggregate preferences/weights together, as when we need to compute the internal inconsistency of an extension (i.e., α in α^{γ}-semantics).

A quantitative study is proposed in [21], where the authors define *Social Abstract Argumentation Frameworks*, which basically associate positive and negative votes to each argument. Afterwards, it is defined how to aggregate these votes together, and how to associate it with an unique social model. This framework has been extend in [5] by considering weights on attacks as well.

An approach to use argumentation as voting methods is instead used in [22]. Here extensions represent a non-conflicting committee to be elected.

In [23] the authors review the works in [19], [24], [13], and [25], focusing on how to relate preference-values and weights, on either arguments or attacks. In

[7], if $R(a, b)$ and $R(b, c)$, a defends c if $W(a, b)$ is worse than $W(b, c)$ (as in [25]), thus the defence is not collective as instead in [2] and this paper, and the attack is not collective as in this work. In [25] the difference between the weight associated with a is related to both the weights of b and c, with the purpose to check how much a defends b (thus obtaining *"varied-strength defeat relations"*).

The two principles in [26] are, *i)* having fewer attackers is better than having more, and *ii)* having more defenders is better than having fewer. The result is the definition of a graded defence $d_{m,n}(\mathcal{E})$, which defines different levels of defence-strength: if $d_{m,n}(\mathcal{E})$ holds, $\mathcal{E}$ is a set of arguments for which each $a \in \mathcal{E}$ does not have at least m attackers that are not counter-attacked by at least n arguments in $\mathcal{E}$. Hence, the notions of defence in [25, 27, 26] follow a different approach and cannot straightforwardly be represented by our framework.

One of the main advantages of the general semiring-based framework proposed here is to be capable of modelling several different WAAFs. This results into a comprehensive computational-model for Weighted Abstract Argumentation. As shown by Theorem 7.3, even classical Dung's semantics can be modelled in the same framework. In the following of this section we report different quantitatively-oriented WAAFs in literature [28, 29, 13], which are all encompassed in our general approach.

For instance, an argument can be seen as a chain of possible events that makes a hypothesis true [28]. The credibility of a hypothesis can then be measured by the total probability that it is supported by arguments. To solve this problem we can use the probabilistic semiring $\langle [0..1], max, \times, 0, 1 \rangle$, where the arithmetic multiplication (i.e., $\times$) is used to compose the probability values together (assuming that the probabilities being composed are independent). In [28] the authors associate probabilities with arguments and defeats. Then, they compute the likelihood of some set of arguments appearing within an arbitrary argument framework induced from the probabilistic framework.

Weights can be also interpreted as subjective beliefs [13]. For example, a weight of $w \in (0, 1]$ on the attack of argument a_1 on argument a_2 might be understood as the belief that (a decision-maker considers) a_2 is false when a_1 is true. This belief could be modelled using probability [13] as well.

The *Fuzzy Argumentation* approach presented in [29] enriches the expressive power of the classical argumentation model by allowing to represent the relative strength of the attack relations between arguments, as well as the degree to which arguments are accepted.

In addition, the weighted semiring $\langle \mathbb{R}^+ \cup \{\infty\}, \min, +, \infty, 0 \rangle$ can model a generic "cost" for the attacks: for example, the number of votes in support of the attack [13], which consequently needs to be minimised.

For what concerns internal inconsistency α, besides [7] we took inspiration from the works in [13], where the authors originally defined the notion of inconsistency budget in an extension. Inconsistency is there considered with the purpose to have a means to compute more (than one, as in [1]) grounded extensions. Attacks are removed from the considered WAAF until the sum of their weights amounts to a value greater than a budget β. The grounded semantics is then computed over all the obtained AAFs (attacks can be removed in different ways). In our approach, the internal inconsistency is used to relax all the classical semantics, not just the grounded one. Our intent is to allow such inconsistency in all of them.

As a reminder, two other notions of weighted defence have been already presented and described in detail in Section 4, i.e., [2] and [3].

11. Conclusions and Future Work

In this work we have defined a new notion of defence for WAAFs. Since defence is collective in the literature (i.e., it considers all the counter-attacks from $\mathcal{B}$ as a whole), our main motivation is to provide a similar view also for all the attacks from a to $\mathcal{B}$, here considered by summing all the attacks weights together. In addition, by casting similar proposals [2, 3] in the same parametric algebraic-framework, it is possible to show all their relations in detail.

Then, we have also shown two different kinds of relaxations of classically crisp concepts in (weighted) Abstract Argumentation. Arguments inside an extension can include conflicts (i.e., they can attack each other), and a new notion of weighted defence (i.e., w-defence) can be relaxed to γ-defence, with the purpose to be less restrictive: γ can "slide" (that is, decrease or increase) with the purpose to model different defences in the literature [2, 3]. Classical inclusions still holds in this new framework, which can also be adopted to encompass Dung's original proposal [1]. An implementation of the framework has been developed as well, and tests show that for small α or γ the average number of extensions slowly increases, thus permitting to catch few very "close" solutions characterised by a low amount of inconsistency.

Inconsistency between the beliefs and/or preferences of agents is ubiquitous in everyday life. By increasing the inconsistency thresholds (both α and γ in our case), we get progressively more solutions: finally, we can prefer solutions obtained with a smaller inconsistency values. This approach permits a finer-grained level of analysis of argument systems than is typically possible [6].

To summarise, we provide a new computational framework for Weighted Abstract Argumentation: we start from a new defence (motivating it), and show how to relax it and how to relax the internal conflict. The benefits of the proposed framework can be summarised by the following points:

- Noise tolerance, through a small internal-threshold α, helps to mitigate the disturbing effects of poorly specified or unsound attacks (*e.g.*, from a *troll*) [5], or errors deriving from the automatic generation of attacks among arguments (*e.g.*, after the process of *argument mining* [17]).
- Relaxing a problem (α and γ) helps to find more or larger solutions. For instance this can be advantageous with the α^γ-stable semantics, since a stable extension [1] is not required to always exist.
- Conversely, constraining a problem can result in a more sceptical approach whose result is a more refined set of extensions, as the case-study in Section 9.2 proves: for instance, stronger (and less) α^γ-stable extensions.

- By defining γ-defence we can soften or harden the impact of weight in the notion of defence: defence may result to be more or less efficacious. This is not obtainable in unweighted frameworks.

- As shown in Section 9, α and γ are independent parameters but influence each other. An agent may prefer to work with extensions either with a better α or a better γ, thus preferring a better internal consistency or a stronger defence.

- The presented framework preserves most of the features of [1], *e.g.*, the implications among semantics are the same ones. For specific values of α and γ (respectively, $\top$ and $\bot$), our framework exactly collapses onto [1].

- We adopt a parametric and general algebraic structure to model all the weights and concepts in this paper. Different proposals (*e.g.*, fuzzy, see Section 10) can be unified by our proposal.

In the future we plan to extend such relaxations to coalitions of arguments [30], and to perform a deeper analysis on real-world cases with weighted networks as in [16]. In addition, as α sums up to the total internal conflict, we plan to count the total external conflict by extending γ-defence to δ-defence, where $\delta = \prod \gamma_a$ (for a any attacker of $\mathcal{B}$).

Moreover, we will investigate the α^γ-grounded semantics, which deserves separate considerations: a straightforward definition, along the line presented in Section 6, would lead to more than one grounded extension (as in [13]). To have a single extension requires a definition alternative to the minimal set-inclusion of α^γ-complete extensions; *e.g.*, we can consider the set of all sceptically accepted arguments (in the α^γ-complete semantics). In the presented framework it is possible to define a single grounded extension that coincides with the intersection of all the α^γ-complete extensions (and with the union of all sceptically-accepted arguments). Under specific conditions, this also corresponds to the minimal α^γ-complete extension (unique in this case). First results in this direction are reported in [31].

We will study how the performance change on different semantics with respect to different graph databases [32, 33]. We will also study two-criteria (α and γ) decision-making procedures to help an agent choose between internal or defence relaxations (as for Figure 22), as introduced at the end of Section 9.

Finally, we will study how the presented framework can be used to model *ranking-based* semantics [34], where a (partial) preference order is defined among the arguments of an AAF. A first step in this direction has been already moved in [35], but our plan is to generalise different ranking-based functions by using our semiring-based framework with weights computed as *Shapely values*.

Acknowledgement

We would like to thank Simone Gabbriellini for the his precious comments and suggestions on the case-study in Section 9.2. This research is partially supported by project "Argumentation 360" (funded by Dept. of Mathematics and Computer Science, University of Perugia), and project "REMIX" (funded by Banca d'Italia and Fondazione Cassa di Risparmio di Perugia).

References

[1] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, Artificial Intelligence 77 (2) (1995) 321–357.

[2] S. Coste-Marquis, S. Konieczny, P. Marquis, M. A. Ouali, Weighted attacks in argumentation frameworks, in: Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR, AAAI Press, 2012, pp. 593–597.

[3] D. C. Martínez, A. J. García, G. R. Simari, An abstract argumentation framework with varied-strength attacks, in: Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, AAAI Press, 2008, pp. 135–144.

[4] S. Bistarelli, U. Montanari, F. Rossi, Semiring-based Constraint Solving and Optimization, Journal of the ACM 44 (2) (March 1997) 201–236.

[5] S. Egilmez, J. Martins, J. Leite, Extending social abstract argumentation with votes on attacks, in: Theory and Applications of Formal Argumentation - Second International Workshop, TAFA, Vol. 8306, Springer, 2013, pp. 16–31.

[6] L. E. Bertossi, A. Hunter, T. Schaub, Introduction to inconsistency tolerance, in: Inconsistency Tolerance [result from a Dagstuhl seminar], Vol. 3300 of Lecture Notes in Computer Science, Springer, 2005, pp. 1–14.

[7] S. Bistarelli, F. Santini, A common computational framework for semiring-based argumentation systems, in: H. Coelho, R. Studer, M. Wooldridge (Eds.), ECAI, Vol. 215 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2010, pp. 131–136.

[8] S. Bistarelli, F. Rossi, F. Santini, A collective defence against grouped attacks for weighted abstract argumentation frameworks, in: Proceedings of the Twenty-Ninth International Florida Artificial Intelligence Research Society Conference, FLAIRS, AAAI Press, 2016, pp. 638–643.

[9] S. Bistarelli, F. Rossi, F. Santini, A relaxation of internal conflict and defence in weighted argumentation frameworks, in: Logics in Artificial Intelligence - 15th European Conference, JELIA, Vol. 10021 of Lecture Notes in Computer Science, Springer, 2016, pp. 127–143.

[10] H. Prakken, G. Vreeswijk, Logics for defeasible argumentation, in: Handbook of philosophical logic, Springer, 2002, pp. 219–318.

[11] S. Bistarelli, F. Gadducci, Enhancing constraints manipulation in semiring-based formalisms, in: ECAI, Vol. 141 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2006, pp. 63–67.

[12] T. S. Blyth, M. F. Janowitz, Residuation theory, Vol. 102, Pergamon press Oxford, 1972.

[13] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, M. Wooldridge, Inconsistency tolerance in weighted argument systems, in: Conf. on Autonomous Agents and Multiagent Systems, IFAAMS, 2009, pp. 851–858.

[14] S. Bistarelli, F. Santini, Conarg: A constraint-based computational framework for argumentation systems, in: IEEE 23rd International Conference on Tools with Artificial Intelligence, ICTAI 2011, IEEE, 2011, pp. 605–612.

[15] P. Erdős, A. Rényi, On the evolution of random graphs, Bull. Inst. Internat. Statist 38 (4) (1961) 343–347.

[16] S. Gabbriellini, F. Santini, A micro study on the evolution of arguments in amazon.com's reviews, in: PRIMA 2015: Principles and Practice of Multi-Agent Systems - 18th International Conference, Vol. 9387 of LNCS, Springer, 2015, pp. 284–300.

[17] A. Peldszus, M. Stede, From argument diagrams to argumentation mining in texts: A survey, International Journal of Cognitive Informatics and Natural Intelligence (IJCINI) 7 (1) (2013) 1–31.

[18] T. J. M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, J. Log. Comput. 13 (3) (2003) 429–448.

[19] L. Amgoud, C. Cayrol, On the acceptability of arguments in preference-based argumentation, in: UAI '98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 1998, pp. 1–7.

[20] S. Modgil, Reasoning about preferences in argumentation frameworks, Artificial Intelligence 173 (9-10) (2009) 901–934.

[21] J. Leite, J. Martins, Social abstract argumentation, in: IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI/AAAI, 2011, pp. 2287–2292.

[22] I. Benedetti, S. Bistarelli, From argumentation frameworks to voting systems and back, Fundamenta Informaticae 150 (1) (2017) 25–48.

[23] C. Cayrol, M. C. Lagasquie-Schiex, From preferences over arguments to preferences over attacks in abstract argumentation: A comparative study, in: 2013 IEEE 25th International Conference on Tools with Artificial Intelligence, IEEE Computer Society, 2013, pp. 588–595.

[24] C. Cayrol, C. Devred, M. C. Lagasquie-Schiex, Acceptability semantics accounting for strength of attacks in argumentation, in: ECAI 2010 - 19th European Conference on Artificial Intelligence, Vol. 215, IOS Press, 2010, pp. 995–996.

[25] S. Kaci, C. Labreuche, Arguing with valued preference relations, in: Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 11th European Conference, ECSQARU, Vol. 6717 of Lecture Notes in Computer Science, Springer, 2011, pp. 62–73.

[26] D. Grossi, S. Modgil, On the graded acceptability of arguments, in: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI, AAAI Press, 2015, pp. 868–874.

[27] L. Amgoud, J. Ben-Naim, D. Doder, S. Vesic, Ranking arguments with compensation-based semantics, in: Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR, AAAI Press, 2016, pp. 12–21.

[28] H. Li, N. Oren, T. J. Norman, Probabilistic argumentation frameworks, in: Theorie and Applications of Formal Argumentation - First International Workshop, TAFA, Vol. 7132 of Lecture Notes in Computer Science, Springer, 2012, pp. 1–16.

[29] M. Schroeder, R. Schweimeier, Fuzzy argumentation for negotiating agents, in: The First International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS, ACM, 2002, pp. 942–943.

[30] S. Bistarelli, F. Santini, Coalitions of arguments: An approach with constraint programming, Fundamenta Informaticae 124 (4) (2013) 383–401.

[31] S. Bistarelli, F. Santini, A Hasse diagram for weighted sceptical semantics with a unique-status grounded semantics, in: Logic Programming and Nonmonotonic Reasoning - 14th International Conference, LPNMR, Vol. 10377 of Lecture Notes in Computer Science, Springer, 2017, pp. 49–56.

[32] S. Bistarelli, F. Rossi, F. Santini, A comparative test on the enumeration of extensions in abstract argumentation, Fundam. Inform. 140 (3-4) (2015) 263–278.

[33] S. Bistarelli, F. Rossi, F. Santini, Not only size but also shape counts: abstract argumentation solvers are benchmark-sensitive, J. Log. Comput. to appear.

[34] E. Bonzon, J. Delobelle, S. Konieczny, N. Maudet, A comparative study of ranking-based semantics for abstract argumentation, in: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI Press, 2016, pp. 914–920.

[35] F. Santini, Graded justification of arguments via internal and external endogenous features, in: Scalable Uncertainty Management - 10th International Conference, SUM, Vol. 9858 of Lecture Notes in Computer Science, Springer, 2016, pp. 352–359.