

Ranking Semantics Based on Subgraphs Analysis

Pierpaolo Dondio

School of Computing, Dublin Institute of Technology
Dublin, Ireland
pierpaolo.dondio@dit.ie

ABSTRACT

We propose a set of novel ranking-based semantics based on a measure of the sensitivity of each argument in an abstract argumentation framework. The sensitivity index is an indicator of how sensitive the label assigned to an argument by an argumentation semantics is, and it is derived from the topology of the graph via a subgraphs analysis coupled with the postulates of the chosen semantics. Using the total rank on arguments induced by such indicator, we propose two ranking-based semantics. We compare the behaviour of our semantics with recent proposals and a widespread set of properties identified in literature. A key feature of our semantics is that the attack relation between arguments keeps the same meaning as found in Dung's abstract semantics. By still relying on Dung's semantics we can soundly deal with any graph configuration, minimize the addition of ad-hoc postulates and provide a clear interpretation of the ranking of arguments.

KEYWORDS

Abstract Argumentation; Ranking-based Semantics.

ACM Reference Format:

Pierpaolo Dondio. 2018. Ranking Semantics Based on Subgraphs Analysis. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10-15, 2018*, IFAAMAS, 9 pages.

1 INTRODUCTION

An abstract argumentation framework [15] consists of a directed graph where nodes represent arguments and arrows represent an attack relation among arguments. A semantics is used to evaluate arguments' acceptability. In the labelling approach [7], this evaluation is done by assigning to each argument a label *in*, *out* or *undec*, meaning that the argument is considered consistently acceptable, non-acceptable or undecided (i.e. no decision can be taken on arguments' acceptability). In Dung's original work, arguments are either fully asserted or not asserted at all, and as a consequence the classical three-value labelling *in*, *undec*, *out* defines an order over the arguments that is often too coarse to support decisions. There is no distinction between arguments with the same label, even if arguments might be in very different positions in the graph.

In order to address this problem, ranking-based semantics were introduced to equip argumentation semantics with a more fine-grained ability to classify arguments. Common to these approaches is the fact that a numerical score is computed for each argument by considering exclusively the topology of the graph. No explicit measure of strength is attached to arguments, and this is the key feature

that differentiates these works from probabilistic ([19][13][22]) or fuzzy argumentation systems ([20][14][12]).

In this paper we propose a topology-based numerical score for arguments called sensitivity index and a set of ranking semantics based on the total rank induced by it.

The idea is to measure how sensitive the label assigned to an argument by a given semantics is by considering how the label changes when the argumentation graph is perturbed. Perturbations are represented by the removal of one or more arguments from the starting graph, modelling the situation in which an argument might be retracted or defeated by the arrival of new evidence. The sensitivity index of argument *a* is computed by analysing the behaviour of the label assigned to *a* over the vertex-induced subgraphs of the complete graph. Informally, the sensitivity index of argument *a* and label *l* is the portion of subgraphs where *a* is labelled *l*. Using the sensitivity index, it is possible to quantify the propensity of an argument to move to a more favourable label, keep the same one or move to a less favourable label if the graph is changed.

We then define a ranking-based semantics based on a combination of Dung's semantics and our sensitivity index. Our ranking-based semantics is first built on the idea of maintaining the coarse ordering identified by the three labels *in*, *undec*, *out*. Then, if two arguments have the same label, the sensitivity index is used to decide how arguments are ranked: the argument with higher propensity of moving to a more favourable label is ranked first.

Our sensitivity index is able to capture minimal changes in the graph. The topology of the graph is fully examined by relying on the subgraphs analysis, and the argumentation semantics are applied without changing their rules or adding extra assumptions. Arguments are still abstract symbolic entities and the notion of *attack* keeps the same meaning as in Dung's abstract argumentation, limiting the need for added postulates.

The paper is organized as follows. The next section describes the required background of abstract argumentation. Section 3 describes our sensitivity index, while section 4 introduces two ranking-based semantics defined using grounded semantics. Section 5 illustrates their properties, while Section 6 describes the extension to preferred semantics. Section 7 contains a discussion of the key features and novelties of our semantics and a comparison with the state of the art. Section 8 concludes the paper.

2 GROUNDED AND PREFERRED SEMANTICS

In this section we describe the required concepts of abstract grounded and preferred semantics introduced by Dung [15].

Definition 1. An argumentation framework *AF* is a pair $\langle Ar, \mathcal{R} \rangle$, where *Ar* is a non-empty finite set whose elements are called arguments and $\mathcal{R} \subseteq Ar \times Ar$ is a binary relation, called the attack relation. If $(a, b) \in \mathcal{R}$ we say that *a* attacks *b*. Two arguments *a*, *b* are rebuttals iff $(a, b) \in \mathcal{R} \wedge (b, a) \in \mathcal{R}$, i.e. they define a symmetric attack. An argument *a* is initial if it is not attacked by any arguments, including itself.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10-15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

An abstract argumentation semantics identifies a set of arguments that can survive the conflicts encoded by the attack relation \mathcal{R} . Dung's semantics require a group of acceptable arguments to be conflict-free (we cannot accept at the same time an argument and its attacker) and admissible (the set of arguments defends itself from external attacks).

Definition 2. (conflict-free). A set $Arg \subseteq Ar$ is conflict-free iff $\forall a, b \in Arg, (a, b) \notin \mathcal{R}$.

Definition 3. (admissible set, complete set). A set $Arg \subseteq Ar$ defends an argument $a \in Ar$ iff $\forall b \in Ar$ such that $(b, a) \in \mathcal{R}, \exists c \in Arg$ such that $(c, b) \in \mathcal{R}$. The set of arguments defended by Arg is denoted $\mathcal{F}(Arg)$. A conflict-free set Arg is *admissible* if $Arg \subseteq \mathcal{F}(Arg)$ and it is *complete* if $Arg = \mathcal{F}(Arg)$.

We follow the labelling approach of [7], where a semantics assigns to each argument a label *in*, *out* or *undec*.

Definition 4. (labelling). Let $AF = (Ar, \mathcal{R})$. A labelling is a total function $\mathcal{L} : Ar \rightarrow \{in, out, undec\}$. We write $in(\mathcal{L})$ for $\{a \in Ar | \mathcal{L}(a) = in\}$, $out(\mathcal{L})$ for $\{a \in Ar | \mathcal{L}(a) = out\}$, and $undec(\mathcal{L})$ for $\{a \in Ar | \mathcal{L}(a) = undec\}$.

Definition 5. (from [7]). Let $AF = (Ar, \mathcal{R})$. A **complete labelling** is a labelling such that for every $a \in Ar$ holds that:

1. if a is labelled *in* then all its attackers are labelled *out*;
2. if a is labelled *out* then it has at least one attacker that is labelled *in*;
3. if a is labelled *undec* then it has at least one attacker labelled *undec* and it does not have an attacker that is labelled *in*.

Definition 6. (grounded and preferred labelling [7]) Given $AF = (Ar, \mathcal{R})$, \mathcal{L} is the grounded labelling iff \mathcal{L} is a complete labelling where $undec(\mathcal{L})$ is maximal (w.r.t. set inclusion) among all complete labellings of AF . \mathcal{L} is the preferred labelling iff \mathcal{L} is a complete labelling where $in(\mathcal{L})$ is maximal (w.r.t. set inclusion) among all complete labellings of AF .




Figure 1: Two Argumentation Graphs G_1 (left) and G_2 (right)

Referring to Figure 1, the grounded labelling assigns the *undec* label to all the arguments of G_1 . Regarding the preferred semantics, there are two complete labellings that maximise the $in(\mathcal{L})$ set: one with $in(\mathcal{L}_1) = \{b\}$, $out(\mathcal{L}_1) = \{a, c\}$, $undec(\mathcal{L}_1) = \emptyset$ and the other with $in(\mathcal{L}_2) = \{a, c\}$, $out(\mathcal{L}_2) = \{b\}$, $undec(\mathcal{L}_2) = \emptyset$. Regarding G_2 , there is only one complete labelling (thus representing both the grounded and preferred labelling), where argument a is *in* (no attackers), b is *out* and c is *in*. Note how a reinstates c .

2.1 Subgraphs notation and labellings

Given $AF = (Ar, \mathcal{R})$ with $|Ar| = n$, and the graph G identified by Ar and \mathcal{R} , $a \in Ar$, we consider the set \mathcal{S} of all the vertex-induced subgraphs of G . We are interested in sets of subgraphs, i.e. elements of $2^{\mathcal{S}}$. We call A and \bar{A} respectively the set of subgraphs where argument a is present and the complementary set of subgraphs where a is absent.

If $Ar = \{a_1, \dots, a_n\}$, a single subgraph s is expressed by an intersection of n sets A_i or (\bar{A}_i) ($i \leq n$) depending on whether the i^{th} argument a_i is or is not contained in s . A set of subgraphs is expressed by combining some of the sets $A_1, \dots, A_n, \bar{A}_1, \dots, \bar{A}_n$ with the connectives $\{\cup, \cap\}$. We write AB to denote $A \cap B$ and $A + B$ for

$A \cup B$. For instance, in Figure 1 left the single subgraph with only b and c present is denoted with $\bar{A}BC$, while AB denotes a set of two subgraphs (ABC and $AB\bar{C}$) where arguments a and b are present and the status of c (not in the expression AB) is not specified.

Given $AF = (Ar, \mathcal{R})$, $a \in Ar$, we define the set of arguments *relevant* to a , called $C_n(a)$, as the set including argument a and all the arguments $x \in Ar$ for which there is a directed path from x to a . For instance, in graph G_1 , $C_n(a) = \{a, b, c\}$, $C_n(b) = C_n(c) = \{b, c\}$.

Based on $C_n(a)$, we define the portion of argumentation framework *relevant* to argument a , called AF_a .

Definition 2.1. Given $AF = (Ar, \mathcal{R})$ and $a \in Ar$, the argumentation framework relevant to a , called AF_a , is an argumentation framework composed only by arguments in $C_n(a)$, that is $AF_a = (C_n(a), \mathcal{R}_a)$, where $\mathcal{R}_a = \mathcal{R} \cap (C_n(a) \times C_n(a))$.

The labelling of a subgraph $s \in \mathcal{S}$ follows the rules of the chosen semantics. For each argument a , A_{IN}, A_{OUT}, A_U are the sets of subgraphs of AF_a where a is present and labelled *in*, *out*, *undec*. Note how in general the three sets are not disjoint. For instance, if we consider graph G_1 and grounded semantics, argument a is labelled *in* in all the subgraphs where a is present and b is not present (and c becomes irrelevant), i.e. $A_{IN} = \bar{A}\bar{B}$ (a set of 2 subgraphs). It is *undec* when all the arguments are present (the single subgraph $A_U = ABC$) while a is *out* when b is present and c is not present, i.e. $A_{OUT} = AB\bar{C}$.

3 SENSITIVITY INDEX OF ARGUMENTS

The sensitivity index of an argument is a quantification of how sensitive the label assigned to each argument is when the argumentation graph is perturbed. We require the sensitivity index to verify the directionality property [3]. This means that the computation of the sensitivity index is entirely based on the set of arguments relevant to a , i.e. $C_n(a)$. We first define the sensitivity index for semantics for which a unique labelling always exists. The idea is to count the portion of subgraphs of AF_a where a has a specific label l . We remind that the set A_l represents the set of subgraphs of AF_a where argument a is present and labelled l by the chosen semantics. The number of subgraphs of AF_a containing argument a is equal to $2^{|C_n(a)|-1}$. For semantics that generate exactly one labelling for each subgraph, the sensitivity index for an argument and label is computed as follows.

Definition 3.1. Let us consider $AF = (Ar, \mathcal{R})$. Given a label $l \in \{i(= in), o(= out), u(= undec)\}$, the sensitivity index of argument $a \in Ar$, label s and semantics x is:

$$S_{x_l}(a) = \frac{|A_l|}{2^{|C_n(a)|-1}} \quad (1)$$

that is the number of subgraphs of AF_a where argument a is labelled l over the total number of subgraphs.

Since $S_{x_l}(a)$ is a proportion, it is $0 \leq S_{x_l}(a) \leq 1$. In this paper, the semantics analysed are grounded and preferred, and therefore $x \in \{g, p\}$.

We call the sensitivity index of an argument the triple $S_x(a) = \langle S_{x_i}(a), S_{x_o}(a), S_{x_u}(a) \rangle$. The following is a straightforward fundamental property for unique status semantics where a labelling always exists.

Property 1. Given $AF = (Ar, \mathcal{R})$, for each $a \in Ar$ it is:

$$S_{x_i}(a) + S_{x_o}(a) + S_{x_u}(a) = 1 \quad (2)$$

3.1 Sample Computations

Initial Argument. If a is an initial argument, we can prove the following lemma:

LEMMA 3.2. Given $AF = \langle Ar, \mathcal{R} \rangle$, for each complete semantics x it holds that $\forall a \in Ar, S_x(a) = \langle 1, 0, 0 \rangle$ if and only if a is initial.

PROOF. If a is initial, then AF_a is reduced to $\langle \{a\}, \emptyset \rangle$ and a is labelled *in* in the only subgraph (that is the complete graph with only a present). Therefore $S_{x_i}(a) = \frac{|A_{in}|}{2^{|A|-1}} = 1$, while $S_{x_u}(a) = S_{x_o}(a) = 0$. If a is not initial, then it is not labelled *in* at least in the subgraph where a is present, one of its attackers b is present (note how a can also attacks itself) and all the other arguments are absent, and therefore $S_{x_i}(a) < 1$. \square

Example. Let us compute the sensitivity index for the two graphs G_1 and G_2 in Figure 1 using grounded semantics. Table 1 summarizes the results.

In G_1 , argument a has $C_n(a) = \{a, b, c\}$. a is labelled *undec* in the full graph. It is labelled *in* when b is removed (2 subgraphs) and *out* when c is removed but b is present. Therefore $S_g(a) = \langle \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$. Arguments b and c are symmetrical rebuttals, and therefore we have $S_g(b) = S_g(c)$. Considering b , it is $C_n(b) = \{c, b\}$. Argument b is *in* when c is removed, and *undec* when c is present. Therefore $S_g(b) = S_g(c) = \langle \frac{1}{2}, \frac{1}{2}, 0 \rangle$.

Regarding graph G_2 , argument a is initial and therefore $S_g(a) = \langle 1, 0, 0 \rangle$. Argument b has two arguments connected and therefore 4 subgraphs to be analysed. Argument b is *in* when both a and c are not present, it is *out* if a is present (2 subgraphs, the presence of c is irrelevant) and it is *undec* when a is absent and c is present. Therefore $S_g(b) = \langle \frac{1}{4}, \frac{1}{4}, \frac{1}{2} \rangle$. Regarding c , it is *in* when b is not present or a is present to defeat b (a total of 3 subgraphs) and it is *undec* when b is present and a is not. Therefore $S_g(c) = \langle \frac{3}{4}, \frac{1}{4}, 0 \rangle$.

Table 1: sensitivity index for graphs G_1 and G_2

Arg	$C_n(a)$	Arg_{in}	Arg_u	Arg_{out}	$S_g(Arg)$
G_1					
a	$\{a, b, c\}$	\overline{AB} (2)	ABC (1)	$AB\overline{C}$ (1)	$\langle \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$
b	$\{b, c\}$	\overline{BC} (1)	BC (1)	\emptyset	$\langle \frac{1}{2}, \frac{1}{2}, 0 \rangle$
c	$\{b, c\}$	\overline{CB} (1)	BC (1)	\emptyset	$\langle \frac{1}{2}, \frac{1}{2}, 0 \rangle$
G_2					
a	$\{a\}$	A (1)	\emptyset	\emptyset	$\langle 1, 0, 0 \rangle$
b	$\{a, b, c\}$	\overline{BAC} (1)	\overline{ABC} (1)	BA (2)	$\langle \frac{1}{4}, \frac{1}{4}, \frac{1}{2} \rangle$
c	$\{a, b, c\}$	$\overline{CB} + ABC$ (3)	\overline{ABC} (1)	\emptyset	$\langle \frac{3}{4}, \frac{1}{4}, 0 \rangle$

Extreme cases. Figure 2 shows the extreme cases for grounded semantics, where S_{g_i} , S_{g_o} and S_{g_u} have the maximum unitary value. We already showed how an initial argument has $S_g(a) = \langle 1, 0, 0 \rangle$. The graph in the middle represents an argument attacked by n unattacked arguments b_i . Argument a is labelled *in* only in the single subgraph where all the n attackers are removed, and therefore $S_{g_i}(a) = \frac{1}{2^n}$ that goes asymptotically to zero for $n \rightarrow \infty$. Since the graph is acyclic, we have $S_{g_o}(a) = 1 - S_{g_i}(a) \approx 1$ for $n \rightarrow \infty$. The graph on the right represents an argument rebutted by n unattacked arguments b_i . Argument a is labelled *in* only in the single subgraph where all the n attackers are removed, and it is labelled *undec* in all the other subgraphs. Therefore $S_{g_u}(a) = 1$ for $n \rightarrow \infty$. Note how also a self-attacking isolated argument has $S_{g_u}(a) = 1$.




Figure 2: Extreme cases: $S_i(a) = 1$, $S_o(a) \approx 1$, $S_u(a) \approx 1$.

4 TWO RANKING-BASED SEMANTICS BASED ON THE SENSITIVITY INDEX

A ranking-based semantics ranks arguments of an argumentation framework. It is defined as follows.

Definition 4.1. Given $AF = \langle Ar, \mathcal{R} \rangle$, a **ranking-based semantics** σ associates to AF a ranking \geq_{AF}^σ on Ar , where \geq_{AF}^σ is a reflexive and transitive relation on Ar . $a \geq_{AF}^\sigma b$ means that a is at least acceptable as b . $a >_{AF}^\sigma b$ is a shortcut for $a \geq_{AF}^\sigma b$ and $b \not\geq_{AF}^\sigma a$.

By ordering the sensitivity index triples $\langle S_{g_i}, S_{g_u}, S_{g_o} \rangle$ of all arguments using the lexicographic order we define a total order on arguments, and therefore we identify a ranking-based semantics called σ_{SLI_x} , where x indicates the semantics used. For instance, in G_1 we obtained that b and c had the same sensitivity index (they have symmetrical positions in the graph) and they ranked higher than a . This is intuitive, since b and c rebut their attackers, while a does not. For the graph G_2 , a had $S_g(a) = \langle 1, 0, 0 \rangle$. Argument b had $S_g(b) = \langle \frac{3}{4}, \frac{1}{4}, 0 \rangle$, lower than a , while c was the lowest-ranked argument with $S_g(c) = \langle \frac{1}{4}, \frac{1}{4}, \frac{1}{2} \rangle$. This is also intuitive: a is initial and therefore ranked first, b is rebutted by c and therefore ranked after a but before c that is rebutted by b but also attacked by a .

However, since the sensitivity index is computed by analysing the subgraphs of an argumentation framework, its value is not linked to the label assigned by the semantics to an argument in the full starting graph. It is possible to have an *out* labelled argument that has a higher sensitivity index than an *in* labelled argument (an example is given in Section 5, Table 2, case 7 versus 14).

Since our aim was to equip Dung's semantics with a more fine-grained additional way of ranking arguments, it is desirable to define a ranking-based semantics that keeps the order defined by the labels $in > undec > out$. Note how this chosen order is an acceptability-focused order and it is only one of the possible orders over labels (see section 7 for further discussion). Therefore we also propose a second ranking-based semantics only partially based on the sensitivity index. In order to rank arguments, we consider the label assigned to each argument by the semantics in the full argumentation graph, and we first order arguments using the labels *in*, *undec*, *out*. This means that an *in* labelled argument is always ranked higher than an *undec* labelled argument and an *out* labelled one. If two arguments have the same label, the sensitivity index is used to decide what argument is ranked first. While Dung's semantics identifies the acceptability status of an argument, the sensitivity index measures the propensity of an argument to move to a more favourable label. We call σ_{LSI_x} this second ranking-based semantics combining label l and sensitivity index. The ranking of σ_{LSI_x} is defined by the lexicographic order of the following 4-tuple $LS_x(a)$, called the *semantic index* of a :

$$LS_x(a) = \langle l, S_{x_i}(a), S_{x_u}(a), S_{x_o}(a) \rangle \quad (3)$$

where l is the label assigned by the semantics to a and $S_{x_i}(a)$, $S_{x_u}(a)$, $S_{x_o}(a)$ are the values of the sensitivity index triple.

4.1 The Grounded case: σ_{LSI_g} and σ_{SI_g} semantics.

The ranking-based semantics σ_{LSI_g} and σ_{SI_g} are identified by the sensitivity index and grounded semantics. Their behaviour is analysed with a series of examples listed in Table 2, covering key situations in abstract argumentation such as attack, reinstatement, rebuttals and accrual. The graphs are ranked by the semantic index $LS_g(a)$. The last column on the right is explained in section 6.

Reinstatement. If an argument a is attacked by b that in turn is attacked by c , c might reinstate a and restore its *in* label. Reinstatement happens in the σ_{LSI_g} semantics, first of all because it happens for Dung's grounded semantics. However, also the sensitivity index of a is improved. If a is attacked by an initial argument (case 14), its sensitivity index is $S_g(a) = \langle \frac{1}{2}, 0, \frac{1}{2} \rangle$. If a third argument reinstates a by unidirectionally defending it (case 5), the value of S_{g_i} is increased from $\frac{1}{2}$ to $\frac{3}{4}$ and S_{g_o} is consequently decreased. However, S_{g_i} is not reinstated to the full value of an initial argument. The reinstatement is therefore always partial. If we increase the length of the odd-length reinstatement chain the value of S_{g_i} decreases (for a chain of 5 arguments S_{g_i} is $\frac{11}{16} < \frac{3}{4}$) and it goes to $\frac{2}{3}$ with an infinite chain of arguments.

The longer the chain, the less the argument is reinstated. In general, arguments closer to a have a higher impact on its sensitivity index. For instance, let us consider argument a attacked by a chain of 4 arguments b, c, d, e (b the nearest). There are $2^4 = 16$ subgraphs to be assigned. When b is removed, a is labelled *in* and three arguments (c, d, e) are disconnected from a and therefore their presence (or absence) in the graph is irrelevant to label a . Therefore a set of $2^3 = 8$ subgraphs (half of the total) is assigned to the set A_{in} in one step, and $\frac{1}{2}$ added to S_{g_i} .

Accrual. In the σ_{LSI_g} and σ_{SI_g} semantics the effect of arguments accrue. Case 1 shows the effect of reinstating argument a with two arguments compared to one (case 5). Argument a has a higher $S_g(a)$ in case 1 than in case 5. The same is for two attacking arguments (case 15) compared to a single attack (case 14), since $S_g(a)$ is lower in case 15.

Attack. Case 15 shows the effect of an attacking argument. The value of S_{g_i} is diminished and transferred to S_{g_o} .

Cases 6 and 7 are interesting. Argument a has a stronger semantics (and sensitivity) index in case 6, where a single argument reinstates a by attacking all the three attackers. The reason is because a in case 6 is slightly more stable than in case 7, since in case 7 the removal of one of the three arguments is enough to change the label of a , while in case 6 there is only one way to change the label of a , that is the removal of the only defender present in the graph. An agent willing to change the label of a would have three ways to do so in case 7 and only one in case 6.

Rebuttals. Two rebuttal arguments are shown in case 8. The semantic index is $\langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle$, and therefore both the semantics and the sensitivity index are higher than argument a unidirectionally attacked by b (case 14). Therefore rebuttal attacks have less impact than unidirectional attacks.

Case 4 is interesting. It is a reinstatement chain where argument a also rebuts its attacker. The fact that argument a rebuts its attacker should place a in case 4 in a better position than case 5. This is the case, and the value of S_{g_o} of case 5 is transferred to S_{g_u} in case 4.

The comparison of case 10 and case 5 is also interesting. In case 10 argument a is reinstated by an argument rebutting the attacker of




Figure 3: Lemma 5.2: partitioning the argumentation graph

a rather than unidirectionally attacking it (as in case 5). Therefore, we expect that the reinstatement on a will be weaker. Indeed, in case 10 the semantics index is $\langle un, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$ as compared to $\langle in, \frac{3}{4}, 0, \frac{1}{4} \rangle$ for case 4. Not only the label moves from *in* to *undec*, but also S_{g_i} is diminished since part of it is now assigned to S_{g_u} .

The grounded semantics strength is not able to distinguish between cases 8 and 9, that is a single rebuttal versus a chain of 3 rebuttals. The reinstatement with rebuttals does not happen since, as it happens with grounded semantics, the extra rebutting argument is not able to label *out* the attacker of a .

Triads. Cases 11 to 13 describe the combinations of three arguments. Case 11 (cycle of 3 arguments) and case 13, known as the *floating assignment*, will be analysed when we introduce the preferred semantics index in section 6. The comparison between cases 8 and 12 also shows that rebuttal attacks accrue since some of the value of S_{g_i} of case 8 is transferred to S_{g_u} in case 12.

5 PROPERTIES OF σ_{SI_g} AND σ_{LSI_g}

We analyse the properties of our two semantics with respect to a subset of properties identified by [5] in their recent survey of ranking-based semantics. We first recall some definitions due to Cayrol [9] and used by [5] in their analysis.

Definition 5.1. Given $AF = \langle A, \mathcal{R} \rangle$, $a \in A$, we define $R^-(a) = \{b \in A \mid R(b, a)\}$ the set of all the **direct attackers** of a . We call **indirect attackers** of a all the arguments for which there is a path $P(b, a)$ from b to a and the length of the path is odd. The **defenders** of a are all the arguments for which there is an even- and non-zero length path from b to a .

We also provide this useful lemma:

LEMMA 5.2. *Let us consider $AF = \langle A, \mathcal{R} \rangle$, argument $a \in A$ and its set of relevant arguments $C_n(a)$. If the graph identified by $C_n(a)$ can be divided into two graphs C_1 and C_2 both containing a and such that $C_2 \cup C_1 = C_n(a)$, $C_2 \cap C_1 = a$ and there are no attacks from C_1 to C_2 and vice-versa except for attacks to a , then*

$$S_{g_i}(a) = S_{g_i}^1(a)S_{g_i}^2(a)$$

$$S_{g_o}(a) = S_{g_o}^1(a) + S_{g_o}^2(a) - S_{g_o}^1(a)S_{g_o}^2(a)$$

$$S_{g_u}(a) = S_{g_u}^1(a)S_{g_u}^2(a) + S_{g_u}^1(a)S_{g_u}^2(a) + S_{g_u}^1(a)S_{g_o}^2(a)$$

where $S_{g_i}^1(a)$ (and respectively $S_{g_i}^2(a)$) is the sensitivity index of argument a considering only the subgraph C_1 (respectively C_2).

Lemma 5.2 can be used when a set of relevant arguments is added to argument a , and the new added arguments and the old ones do not interact. Lemma 5.2 guarantees that the semantic index of argument a is always decreased by the new addition.

A set of properties analysed in this section is based on the definition 5.1 of attackers and defenders. However, the definition of attackers and defenders of argument a does not distinguish between arguments for which there are multiple paths to a , some of them odd-length and some even-length. Moreover, being an indirect attacker or defender does not reveal how far the argument is from a . Finally, both a rebuttal of a and a unidirectional attacker of a are both direct attackers, but they should not be treated in the same

Table 2: The Behaviour of σ_{LSI_g} and σ_{LSI_p} . In case of acyclic graph, the label *in* means in_{sk} as well

R	AF	Comment	σ_{LSI_g}	σ_{LSI_p}
1		Two reinstating arguments	$\langle in, \frac{7}{8}, 0, \frac{1}{8} \rangle$	
2		Chain of n arguments	$\langle in, \sum_{k=0}^{2k+1 \leq n} \frac{1}{2^{2k+1}}, 0, 1 - \sum_{k=0}^{2k+1 \leq n} \frac{1}{2^{2k+1}} \rangle$	
3		Chain of 5 arguments	$\langle in, \frac{11}{16}, 0, \frac{5}{16} \rangle$	
4		Reinstatement and rebuttal attack	$\langle in, \frac{3}{4}, \frac{1}{4}, 0 \rangle$	$\langle in_{sk}, \frac{7}{8}, 0, \frac{1}{8} \rangle$
5		Reinstatement	$\langle in, \frac{3}{4}, 0, \frac{1}{4} \rangle$	
6		A single defender and multiple attacks	$\langle in, \frac{9}{16}, 0, \frac{7}{16} \rangle$	
7		Multiple reinstatement chains	$\langle in, \frac{27}{64}, 0, \frac{37}{64} \rangle$	
8		Rebutting arguments	$\langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle$	$\langle in_{sk}, \frac{3}{4}, 0, \frac{1}{4} \rangle$
9		Chain of rebuttals arguments	$\langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle$	$\langle in_{cr}, \frac{3}{4}, 0, \frac{1}{4} \rangle$
10		Reinstatement with rebuttal attack	$\langle un, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$	$\langle in_{cr}, \frac{5}{8}, 0, \frac{3}{8} \rangle$
11		Cycle of three attacking arguments	$\langle un, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$	$\langle un_{sk}, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$
12		Two rebuttal attacks	$\langle un, \frac{1}{4}, \frac{3}{4}, 0 \rangle$	$\langle in_{cr}, \frac{5}{8}, 0, \frac{3}{8} \rangle$
13		Floating assignment	$\langle un, \frac{1}{4}, \frac{1}{4}, \frac{1}{2} \rangle$	$\langle out_{sk}, \frac{1}{4}, 0, \frac{3}{4} \rangle$
14		Single attack	$\langle out, \frac{1}{2}, 0, \frac{1}{2} \rangle$	
15		Accrual of attacks	$\langle out, \frac{1}{4}, 0, \frac{3}{4} \rangle$	$\langle out_{sk}, \frac{1}{4}, 0, \frac{3}{4} \rangle$

way, since they might have different impact on a . For these reasons, the majority of the properties based on direct/indirect attackers are not generally valid for our semantics, but more specific versions are. Table 3 shows the properties considered and which ones are satisfied by Dung’s grounded semantics (column *Gr*), the sensitivity index alone (σ_{SI_g}) and the semantic index (σ_{LSI_g}).

The **abstraction property (Abs)** requires the ranking to be fully defined from the graph topology. Both σ_{SI_g} and σ_{LSI_g} satisfy it, since both grounded semantics and our sensitivity index do.

The **independence property (Ind)** states that the ranking of two arguments a and b should be independent of any argument that is neither connected to a nor b . Again, this is true for both grounded semantics and our definition of sensitivity index (and therefore for the semantics index as well).

The **void preference (VP)** property states that an unattacked argument is ranked strictly before an attacked one. Both our semantics satisfy this, since an argument has the highest score if and only if it is initial (lemma 3.2).

Our semantics satisfy both the **total (Tot)** property (since all the arguments are ranked), and the **non-attacked equivalence (NaE)**, since all the initial arguments are ranked the same.

Table 3: Properties of our Ranking-based semantics

Property	Gr	σ_{SI_g}	σ_{LSI_g}	Property	Gr	σ_{SI_g}	σ_{LSI_g}
Abs	✓	✓	✓	Ind	✓	✓	✓
VP	×	✓	✓	DP	×	×	×
Tot	✓	✓	✓	NaE	✓	✓	✓
AvsFD	✓	×	✓	SC	×	×	×
CP	×	×	×	QP	✓	×	×
⊕DB	×	×	×				
+AB	×	✓	✓	+DB	×	×	×
↑AB	×	✓(*)	✓(*)	↑DB	×	✓(*)	✓(*)
CT	×	×	×	SCT	×	×	×

(*) Only if \geq is used instead of $>$

The **cardinality precedence (CP)** states that if $|R^-(a)| > |R^-(b)|$ then $b > a$. The property is not verified by our two semantics, since the semantics index of a depends on the semantics index of each attacker and not on the number of attackers. The **defence precedence (DP)** states that if $|R^-(a)| = |R^-(b)|$ but a has at least one defender while b has none, then $a > b$. The property is not verified, since b might be without defenders but attacked

by rebuttal arguments as in the graph in Figure 5. Here a and b have the same number of direct attackers and a is defended, but $LS_g(b) = \langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle > LS_g(a) = \langle un, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$. The **self-contradiction (SC)** property states that a self-attacking argument should be ranked lower than a non-self-attacking one, and it is not verified by any of our semantics since, for instance, an isolated self-attacking argument is ranked higher than an argument attacked by an initial argument.

The **quality precedence (QP)** property considers the rank of the direct attackers, and it is formalized as follows:

$$\exists c \in R^-(b).s.t.\forall d \in R^-(a), c > d \implies a > b \quad (4)$$

The property states that if b is attacked by c and c is ranked higher than all the direct attackers d of a , then a is ranked higher than b . This is verified by grounded semantics, but it is not verified by our two semantics. Indeed a direct attacker c can be stronger than d , but c could be a rebuttal and d a unidirectional attack, so that the effect of d is stronger than c . This is exemplified in Figure 4, where a and b are both labelled *undec*, and $LS_g(c) = \langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle > LS_g(d) = \langle un, \frac{1}{4}, \frac{3}{4}, 0 \rangle$ but $LS_g(b) = \langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle > LS_g(a) = \langle un, \frac{1}{2}, \frac{3}{8}, \frac{1}{8} \rangle$.




Figure 4: A counter-example for Quality Precedence

A **defence branch** of argument a is an even-length chain of arguments ending with a , while an **attack branch** is an odd-length chain of arguments ending with a .

The **strict addition of a defence branch (\oplus DB)** states that adding a defence branch to an argument improves its ranking. This is not verified for lemma 5.2 ($S_{g_i}(a)$ strictly decreases). The more specific **addition of a defence branch (+DB)** states that adding a defence branch to any attacked argument improves its ranking. This property is not verified, again for lemma 5.2. The same lemma proves the **addition of an attack branch (+AB)** property, that states that adding an attack branch will decrease the ranking of a .

The \uparrow **AB** property states that **increasing the length of an attack branch** (that is not also a defence branch) of a will increase the ranking of a . If we call b the argument ending the attack branch before the addition of new arguments, for lemma 5.2 b will decrease its ranking and consequently a could have its ranking improved, unchanged (for instance, in case of a chain of rebuttal attacks) but not decreased. So the property is verified if we use \geq rather than $>$. The same argument is used to prove that **increasing the length of a defence branch** (that is not an attack branch) of a decreases the ranking of a (property \uparrow **DB**).

The **counter-transitivity (CT)** property states that if the attackers of b are as numerous and acceptable (=of higher rank) than the attackers of a , then $a \geq b$. In the **strict counter-transitivity (SCT)** property only the direct attackers are considered. None of the properties is verified, again for the presence of potential rebuttal attacks. The graph in Figure 5 exemplifies this. Arguments a and b have the same number of direct attackers with the same semantic strength. However, $LS_g(b) = \langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle > LS_g(a) = \langle un, \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$.

The **attack vs. full defence (AvsFD)** property states that in an acyclic graph an argument a without any attack branch is ranked higher than an argument b only attacked by one non-attacked argument. This is verified by σ_{LSI_g} since a is labelled *in* and b *out*,




Figure 5: A counter-example for properties DP and CT

but not for σ_{SI_g} since the sensitivity index does not satisfy it. For instance argument a in Table 2 case 14 (a attacked by an initial argument) has a higher S_g than case 7 (a fully defended).

6 EXTENSION TO PREFERRED SEMANTICS

A preferred labelling is a complete labelling that maximizes the set of *in* labelled arguments. It might produce more than one labelling for the same argumentation graph, and therefore an argument a might be labelled in different ways in the same subgraph. In abstract argumentation we distinguish between a label l assigned *sceptically* ($=a$ is labelled l in all the valid labellings) and *credulously* ($=a$ is labelled l in at least one valid labelling).

The question is therefore how to count subgraphs in the computation of the preferred sensitivity index S_p in presence of multiple labellings. Two proposals are possible. Given a label l and an argument a , we could use the *credulous* assignment to count the subgraphs where a has a given label. Under this proposal, the same subgraph s could be counted for more than one label if a has different labels in different labellings of s , with the consequence that property 1 is no more guaranteed and in general $S_{p_i}(a) + S_{p_u}(a) + S_{p_o}(a) \geq 1$.

If we use the *sceptical* label assignment, an argument could not have a sceptically assigned label in some subgraphs and therefore these subgraphs will not be counted in the computation, and $S_{p_i}(a) + S_{p_u}(a) + S_{p_o}(a) \leq 1$.

Our proposal applies the insufficient reason principle to compute a value for $S_p(a)$ that still satisfies property 1. Let us presume that there are n valid labellings for the same subgraph s , and in n_{in} labellings argument a is labelled *in*, in n_{un} labellings is *undec* and in n_{out} labellings is *out*. It is $n = n_{in} + n_{out} + n_{un}$. Since all the labellings are equal for the semantics, the insufficient reason principle suggests to equally split the contribution of the subgraph s between $S_{p_i}(a)$, $S_{p_o}(a)$, $S_{p_u}(a)$ proportionally to the number of labellings where a is respectively labelled *in*, *out*, *undec*. For instance, the subgraph will contribute to $S_{p_i}(a)$ with a value of $\frac{n_{in}}{n}$.

The definition of the sensitivity and semantics index is modified as follows. Given $AF = \langle A, \mathcal{R} \rangle$, $a \in A$, we first define a function $P(s, l, a) : \mathcal{S} \times \{in, un, out\} \times \mathcal{A} \rightarrow \mathbb{R}$ that for each subgraph $s \in \mathcal{S}$, label l and argument a returns the quantity $\frac{n_l}{n}$, the proportion of labellings contributing to the label l for argument a in the subgraph s . The sensitivity index for preferred semantics is:

$$S_p(a) = \frac{\sum_{s \in \mathcal{S}} P(s, l, a)}{2^{|C_n(a)|-1}} \quad (5)$$

Note how this definition is also valid for unique status semantics, where $P(s, l, a)$ will be either one or zero.

As we did for the grounded case, we define two ranking-based semantics σ_{SI_p} and σ_{LSI_p} for the preferred case. The *preferred* sensitivity index S_p is used to define the semantics of σ_{SI_p} , while σ_{LSI_p} is defined using the preferred semantics index LSI_p .

The semantics index for preferred semantics has still the same shape $\langle l, S_i, S_u, S_o \rangle$ but now the label l has the form l_{acc} , where $l \in \{in, undec, out\}$ and $acc \in \{sk, cr\}$ for sceptical and credulous label assignment. Arguments are ranked by their label $in > undec > out$, but the sceptical assignments are now ranked higher than the

credulous ones. Since an argument can be credulously accepted and rejected in the same subgraph, by definition we consider the most favourable label for each argument.

If two arguments have the same label l_{acc} , the preferred sensitivity index $S_p(a)$ is used to rank arguments.

6.1 The behaviour of σ_{SI_p} and σ_{LSI_p} semantics

If the graph is acyclic σ_{SI_p} and σ_{LSI_p} coincide with σ_{SI_g} and σ_{LSI_g} . For cyclic graphs, the preferred values might differ significantly.

Rebuttals. In case of two rebuttal arguments a and b , there are two subgraphs to be considered for the computation of $S_p(a)$ and $LS_p(a)$. In the subgraph s_1 with only argument a present, a is labelled *in*, while in the complete graph s_2 there are two symmetrical labellings: l_1 where a is *in* and l_2 where a is *out*. Therefore it is $P(s_2, in, a) = P(s_2, out, a) = \frac{1}{2}$ and *half* of the subgraph is assigned to S_{p_i} and *half* to S_{p_o} .

The Semantics index is therefore $LS_p(a) = \langle in_{cr}, \frac{3}{4}, 0, \frac{1}{4} \rangle$. In the grounded case it was $LS_g(a) = \langle un, \frac{1}{2}, \frac{1}{2}, 0 \rangle$, and therefore $S_{g_u} = \frac{1}{2}$ is equally split between S_{p_i} and S_{p_o} .

Note also how a rebutting argument under preferred semantics has a semantics strength $\langle in_{cr}, \frac{3}{4}, 0, \frac{1}{4} \rangle$ similar to a reinstated argument $\langle in_{sk}, \frac{3}{4}, 0, \frac{1}{4} \rangle$, with the only difference of the sceptical assignment. This is in line with the behaviour of preferred semantics where an argument can defend (=reinstate) itself by rebutting attacking arguments.

Graph G_1 (Figure 1). The result of the computation of LS_p for graph G_1 is displayed in Table 4. Note how in each subgraph s where the rebutting cycle of argument b and c is present, the subgraph is equally split between the *in* and *out* label since it is $P(s, in, b) = P(s, out, b) = P(s, in, c) = P(s, out, c) = \frac{1}{2}$ (this is the meaning of the $\frac{1}{2}$ beside some of the subgraphs in Table 4). The ranking induced by σ_{LSI_p} agrees with σ_{LSI_g} , since arguments b and c are ranked before a .

Table 4: Preferred sensitivity index for graph G_1

<i>Arg</i>	<i>Arg_{in}</i>	<i>Arg_{un}</i>	<i>Arg_{out}</i>	$LS_p(Arg)$
a	$\overline{AB} (2) + ABC (\frac{1}{2})$	\emptyset	$ABC (\frac{1}{2}) + \overline{ABC} (1)$	$\langle in_{cr}, \frac{5}{8}, 0, \frac{3}{8} \rangle$
b	$\overline{BC} (1) + BC (\frac{1}{2})$	\emptyset	$BC (\frac{1}{2})$	$\langle in_{cr}, \frac{3}{4}, 0, \frac{1}{4} \rangle$
c	$\overline{CB} (1) + BC (\frac{1}{2})$	\emptyset	$BC (\frac{1}{2})$	$\langle in_{cr}, \frac{3}{4}, 0, \frac{1}{4} \rangle$

We now reconsider Table 2 examples. The last column displays the value of σ_{LSI_p} for each case. Both of the semantics are not able to distinguish between case 8 (a rebutting b) and case 9 (a chain of 3 rebuttals a, b, c). The reinstatement with rebuttal attacks does not happen. Under grounded semantics, this is because the third rebuttal c is not able to defeat b and reinstate a , under preferred semantics because a and b are able to defend themselves.

Cases 11-13 describe the combinations of three arguments. As it happens with Dung's semantics, there is no difference between the grounded and preferred index for case 11 (cycle of three arguments). Case 13 is known as the floating assignment example. Here argument a is sceptically labelled *out*, even if its attackers are credulously accepted (but, in turn, they defeat a in each labelling). The example shows how both the grounded sensitivity and semantic index can have a value greater than the preferred counterparts.




Figure 6: a and b are ranked differently by σ_{LSI_g} and σ_{LSI_p}

6.2 Properties

The following is a simple property linking $S_g(a)$ and $S_p(a)$.

LEMMA 6.1. *For every argument it holds that: $S_{g_i} \leq S_{p_i}, S_{g_o} \leq S_{p_o}, S_{g_u} \geq S_{p_u}$.*

The proof is straightforward. Every time grounded semantics assigns a label *in* or *out*, the same label is also assigned by the preferred semantics (since we are dealing with an acyclic portion of the graph), while when the grounded semantics assigns the label *undec*, the preferred semantics can assign any label sceptically or credulously (see for instance the floating assignment example).

We also checked if there is a bijective function between σ_{LSI_g} and σ_{LSI_p} , that is if the two rankings always agree. The answer is negative, since it is possible to build cases where σ_{LSI_g} and σ_{LSI_p} differ. Referring to Figure 6, we consider argument a and b . We have, $S_g(a) = \langle un, \frac{1}{4}, \frac{1}{4}, \frac{1}{2} \rangle > S_g(b) = \langle out, \frac{1}{4}, \frac{1}{8}, \frac{5}{8} \rangle$ while in the preferred case it is $S_p(a) = \langle out_{sk}, \frac{1}{4}, 0, \frac{3}{4} \rangle < S_p(b) = \langle out_{sk}, \frac{5}{16}, 0, \frac{11}{16} \rangle$. The example shows that σ_{SI_g} and σ_{SI_p} differ as well.

7 DISCUSSION

In this section we discuss the features and the novelty of our semantics with reference to the state-of-the-art.

Keeping Dung's semantics and minimizing ad-hoc postulates. The computation of a ranking-based semantics is a combination of topology analysis and a set of ad-hoc postulates. Our approach is not different. However, we claim that our semantics keep the postulates to a minimum by relying on an unchanged version of Dung's semantics. Arguments are symbolic entities and we do not change the original meaning of the attack relation. Each subgraph is an argumentation framework where Dung's semantics are applied. By relying on Dung's semantics, not only we minimize the introduction of ad-hoc postulates, but we can soundly deal with any graph configuration.

A clear interpretation. We also claim that our ranking semantics has a clear meaning and an intuitive interpretation. The sensitivity index has a clear definition: it is the proportion of subgraphs where an argument label holds.

Our ranking has also an intuitive meaning. Given two arguments with the same label, the stronger argument is the one that has a higher chance to move to a more favourable label when the graph configuration changes. Since arguments are defeasible, the removal of an argument represents a common dynamic situation where arguments are retracted or defeated by new arguments.

Interpretation and ad-hoc postulates in the state-of-the-art. In all the ranking-based semantics arguments are still symbolic entities. However, attacks have the effect of *diminishing* an argument acceptability status. Postulates are added to capture intuitions about how arguments should diminish their acceptance given a graph topology. Postulates model this diminishing effect by considering the number of direct and/or indirect attackers, their *strength* and by defining functions to aggregate and combine arguments (see for instance [1][24][2][4][18][21][17]).

For instance, [24] quantifies the diminishing effect by considering the number of direct attackers, [1] compares arguments by counting the number of paths ending to them, [23] assigns strength to arguments based on the results of the equilibrium of a two-person game, [21] computes an argument score as the solution to a system of equations, where each equation describes supporters and attackers of each argument.

However, the nature of what is diminished is not clear, since arguments are still symbolic entities with no *strength* attached. Moreover, it could be hard to fit sensible postulates to the potentially intricate shape of an argumentation framework and to keep a clear interpretation of the numerical scores inducing the ranking. The score is somehow described as a *degree of acceptability*, a proxy for the *strength* of an argument. However, it is problematic to interpret these ad-hoc numerical scores as a proxy for the strength of arguments resulting from the argumentation. None of these numerical scores can be somehow mapped to other quantitative approaches to reasoning, such as degrees of truth or probabilities. Their non-probabilistic nature limits their usefulness in decision-making, since it is possible to build a *Dutch Book argument* [10] [25] against an agent deciding on the base of those numerical scores. What can be safely said about these numerical scores is that they quantify a topology-based property. Our sensitivity index is not different, however it is fully based on Dung's semantics and it has also a probabilistic interpretation.

Probabilistic interpretation. The sensitivity index quantities S_x are probability measures. Indeed, they represent the portion of subgraphs where a has a specific label. Given an argument a , if we consider the probabilistic argumentation framework defined over the relevant graph AF_a , where $p(a) = 1$, every other argument has the same probability $\frac{1}{|C_n(a)|-1}$ and all the arguments are independent, then S_x is the probability that a is labelled x , as computed by the *constellation approach* [22]. $p(a) = 1$ is set to guarantee that argument a is present in all the subgraphs of the probabilistic argumentation framework. Note how the probability distribution over arguments in $C_n(a)$ is the maximum entropy distribution, and reflects the fact that we only know the topology of the graph and all arguments are equal.

Properties. Our ranking-based semantics handle challenging cases without assuming complicated axioms by still relying on Dung's abstract semantics. We consider again the properties listed in section 5 and we argue that some of them fail to capture the complexity of an argumentation framework, resulting too specific and with little justification.

The properties working with attackers and defenders fail to consider the presence of rebuttal attacks, the distance between the attacking and the attacked argument a and the potential simultaneous presence of odd- and even-length paths from attacking arguments to a . The properties regarding the addition of a defence or attack branch (+AB +DB), are in our opinion counter-intuitive. By adding a new set of arguments to a , a direct attacker b is always added. Independently from the length of the branch added, b is not fully defeated (since there is no full reinstatement in any ranking-based semantics analysed) and therefore b will have an impact on a , potentially decreasing its ranking. The properties regarding the increase of a branch fail to discriminate between rebuttal attacks and the presence of odd and even paths from b to a . The *cardinality precedence* is a naive local property not considering the quality of

the attacks, but the *quality of the attacks* property, as stated, fails again to consider rebuttal attacks.

Computational issues. The computation of our indices shares the same challenges as probabilistic argumentation. It has above-polynomial complexity, but it can benefit from studies in that area, including approximation and optimization algorithms by [16], the recursive computation by [13] and the subgraph-based analysis by [11]. In [13], the set A_l needed to compute our sensitivity index is found by traversing the transpose graph (a graph with reversed arrows) from a down to its attackers, propagating the constraints of the grounded labelling. A merit of using this recursive algorithm is that the values of our sensitivity index would be seen as *propagating* over the argumentation graph from attackers to attacked arguments, rather than be the result of an analysis of unrelated subgraphs as the definition suggests.

Variations. A first variation is considering the removal of individual attack links as well, producing an even more fine-grained ranking. Another variation would be to consider other kinds of graph perturbation beyond the removal of argument. For instance, rather than be removed (equivalent to be *out*-labelled), an argument could be made undecided, modelling the situation where a rebuttal argument is added to it.

A deeper conceptual variation concerns the order induced by the labels assigned to arguments. We have assumed the $in > undec > out$ order. This is an acceptability-focused order that is just one of the possible orders. In an information-focused context, the fact that an argument is labelled *in* or *out* is equally important, while the label *undec* is superseded by the two other labels. Even if the generic framework proposed in this paper is abstract enough to accommodate this alternative ranking, the properties of the resulting ranking might vary substantially. Moreover, other labellings beyond the classical three values labelling have been proposed ([6]) and could be explored. Finally, Dung's framework has been extended in numerous ways, for instance by introducing attacks on attacks or the relation of support between arguments ([8]). The generalization of our framework to these extensions is a challenge for future works, including the possibility of dealing with partial orders on arguments rather than total orders.

8 CONCLUSIONS

In this paper we proposed a set of ranking semantics based on the concept of sensitivity index. The index is an indicator of how sensitive a label assigned to an argument by an argumentation semantics is. This numerical indicator is derived from the topology of the graph via a subgraph analysis, coupled with the postulates of the chosen semantics. A key feature of our ranking-based semantics is that the attack relation between arguments keeps the same meaning as found in Dung's semantics. By still relying on Dung's semantics, we can soundly deal with any graph configuration without additional ad-hoc postulates. We believe the numerous examples illustrated, the comparison with recent proposals and a widespread set of properties identified in literature have shown the soundness of the ranking produced by our semantics.

REFERENCES

- [1] Leila Amgoud and Jonathan Ben-Naim. 2013. Ranking-based semantics for argumentation frameworks. In *International Conference on Scalable Uncertainty Management*. Springer, 134–147.
- [2] Leila Amgoud, Jonathan Ben-Naim, Dragan Doder, and Srdjan Vesic. 2016. Ranking Arguments With Compensation-Based Semantics.. In *KR*. 12–21.

- [3] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26, 4 (2011), 365–410.
- [4] Philippe Besnard and Anthony Hunter. 2001. A logic-based theory of deductive arguments. *Artificial Intelligence* 128, 1-2 (2001), 203–235.
- [5] Elise Bonzon, Jérôme Delobelle, Sébastien Konieczny, and Nicolas Maudet. 2016. A Comparative Study of Ranking-Based Semantics for Abstract Argumentation.. In *AAAI*. 914–920.
- [6] Gerhard Brewka, Sylwia Polberg, and Stefan Woltran. 2014. Generalizations of Dung frameworks and their role in formal argumentation. *IEEE Intelligent Systems* 29, 1 (2014), 30–38.
- [7] Martin WA Caminada and Dov M Gabbay. 2009. A logical account of formal argumentation. *Studia Logica* 93, 2 (2009), 109–145.
- [8] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. 2005. Gradual Valuation for Bipolar Argumentation Frameworks.. In *ECSQARU*, Vol. 3571. Springer, 366–377.
- [9] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. 2005. Graduality in Argumentation. *J. Artif. Intell. Res.(JAIR)* 23 (2005), 245–297.
- [10] Bruno De Finetti. 2017. *Theory of probability: A critical introductory treatment*. Vol. 6. John Wiley & Sons.
- [11] Pierpaolo Dondio. 2013. Computing the grounded semantics in all the sub-graphs of an argumentation framework: an empirical evaluation. In *International Workshop on Computational Logic in Multi-Agent Systems*. Springer, 119–137.
- [12] Pierpaolo Dondio. 2014. Multi-Valued and Probabilistic Argumentation Frameworks.. In *COMMA*. 253–260.
- [13] Pierpaolo Dondio. 2014. Toward a computational analysis of probabilistic argumentation frameworks. *Cybernetics and systems* 45, 3 (2014), 254–278.
- [14] Pierpaolo Dondio. 2017. Propagating degrees of truth on an argumentation framework: an abstract account of fuzzy argumentation. In *Proceedings of the Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, April 3-7, 2017*. 995–1002. <https://doi.org/10.1145/3019612.3019717>
- [15] Phan Minh Dung. 1995. 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.
- [16] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. 2015. On the complexity of probabilistic abstract argumentation frameworks. *ACM Transactions on Computational Logic (TOCL)* 16, 3 (2015), 22.
- [17] Dov M Gabbay. 2012. Equational approach to argumentation networks. *Argument & Computation* 3, 2-3 (2012), 87–142.
- [18] Davide Grossi and Sanjay Modgil. 2015. On the Graded Acceptability of Arguments.. In *IJCAI*. 868–874.
- [19] Anthony Hunter. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54, 1 (2013), 47–81.
- [20] Jeroen Janssen, Martine De Cock, and Dirk Vermeir. 2008. Fuzzy argumentation frameworks. In *Information Processing and Management of Uncertainty in Knowledge-based Systems*. 513–520.
- [21] Joao Leite and Joao Martins. 2011. Social abstract argumentation. In *IJCAI*, Vol. 11. 2287–2292.
- [22] Hengfei Li, Nir Oren, and Timothy J Norman. 2011. Probabilistic Argumentation Frameworks. *TAFIA* 7132 (2011), 1–16.
- [23] Paul-Amaury Matt and Francesca Toni. 2008. A Game-Theoretic Measure of Argument Strength for Abstract Argumentation.. In *JELIA*, Vol. 5293. Springer, 285–297.
- [24] Fuan Pu, Jian Luo, Yulai Zhang, and Guiming Luo. 2014. Argument ranking with categoriser function. In *International Conference on Knowledge Science, Engineering and Management*. Springer, 290–301.
- [25] Frank P Ramsey. 1931. Truth and probability (1926). *The foundations of mathematics and other logical essays* (1931), 156–198.