

Convergent Semantics for Weighted Bipolar Argumentation

Zongshun Wang, Yuping Shen

Institute of Logic and Cognition

Department of Philosophy

Sun Yat-sen University, P.R. China

wangzsh27@mail.sysu.edu.cn, shyping@mail.sysu.edu.cn

Abstract

Establishing *convergent semantics* for *weighted argumentation graphs* is a long-standing fundamental issue. Particularly, it is challenging to develop convergent semantics for *weighted bipolar argumentation graphs* (wBAG), which include both *support* and *attack* relations on weighted arguments. Existing semantics in the literature are *not* general enough in the sense that they only apply to *acyclic* graphs or special *cyclic* cases. In this paper, we provide an elegant solution to this issue by adopting the so-called *bilateral gradual semantics*, so that the strength of arguments can be defined as the *limits* of iterative functions that *always* converge for *any* wBAG including cyclic ones. A preliminary experimental analysis shows that our semantics appear quite efficient in calculating argument strength. Overall, this paper offers a solid and promising foundation for weighted bipolar argumentation in theoretical and practical aspects.

Introduction

An *argumentation graph* is a computational model for reasoning and decision-making in complex environments (Dung 1995; Amgoud and Prade 2009; Atkinson et al. 2017). In recent years, establishing *convergent semantics* has emerged as a compelling and influential approach for evaluating weighted arguments (Gabbay and Rodrigues 2015; Amgoud, Doder, and Vesic 2022; Besnard and Hunter 2001; Prakken 2024). The basic idea is to introduce *iterative functions* that take graphs as inputs and produce sequences of values that eventually *converge*. The final value assigned to arguments is called the *strength* or *acceptability degree*. This approach has numerous applications such as belief revision (da Costa Pereira, Tettamanzi, and Villata 2011), social network (Leite and Martins 2011), explainability (Potyka 2021), etc.

Weighted bipolar argumentation graphs (wBAG) consider both *support* and *attack* relations on weighted arguments. Such bipolar graphs have received extensive attention due to their versatile expressive power (Amgoud et al. 2008; Cayrol and Lagasquie-Schiex 2005; Polberg and Hunter 2018), with applications spanning engineering design (Baroni et al. 2015), polling (Rago and Toni 2017),

review aggregation (Rago et al. 2025), etc. However, developing convergent semantics for wBAG has long been challenging. Most semantics in the literature *fail* to converge in *cyclic* graphs (Evripidou and Toni 2014; Baroni et al. 2015; Rago et al. 2016; Amgoud and Ben-Naim 2018b; Potyka 2021; Doder, Vesic, and Croitoru 2021; Rago et al. 2024). The only exception (Mossakowski and Neuhaus 2018) introduces a restricted semantics that considers only the strongest supporter and attacker, which is less *open-minded* (Potyka 2019a,b). As pointed out in (Amgoud and Ben-Naim 2018b; Potyka 2018; Potyka and Booth 2024b; Yin, Potyka, and Toni 2024), an urgent task is to define semantics that are capable of dealing with *any* typology of graphs including cyclic ones.

To address this issue, we adopt the so-called *bilateral gradual semantics* (Wang and Shen 2024) for wBAG, where each argument is assigned two degrees—*acceptability* and *rejectability*—to capture both the positive and negative strength, reflecting a common cognitive evaluation process (Cacioppo, Gardner, and Berntson 1997). Naturally, the support and attack relations are *separately measured* through the acceptability and rejectability degrees in wBAG. This philosophy robustly leads to convergent semantics through iterative functions, laying a solid foundation for well-behaved wBAG semantics.

Towards this end, we first investigate the *principles* for well-behaved wBAG semantics. We incorporate widely-adopted principles from unipolar support graphs (Amgoud and Ben-Naim 2016), unipolar attack graphs (Amgoud et al. 2017), and bipolar graphs (Amgoud and Ben-Naim 2018b). Particularly, we include three notable principles—*Quality Precedence*, *Cardinality Precedence* and *Compensation*—which are known for developing semantics for unipolar graphs, yet to the best of our knowledge, have never been studied for bipolar graphs. Accordingly, we propose three convergent semantics, each defined as the limit of iterative functions that always converge for any wBAG including cyclic ones. Each semantics corresponds to one of the above three notable principles and satisfies most widely adopted principles. A preliminary experimental analysis shows that the performance of our semantics appears quite efficient concerning the number of iterations and the running time for calculating argument strength.

This paper is organized as follows. We first introduce

basic concepts concerning wBAG and bilateral gradual semantics. Next, we illustrate desirable principles for wBAG semantics and study their interrelation. Then we present our main contribution, i.e., three wBAG semantics that converge for arbitrary graphs including cyclic ones. Satisfied principles are summarized in Table 2 and preliminary experimental results are depicted in Figure 1. The paper ends with a discussion and conclusion.

Preliminaries

A *weighted bipolar argumentation graph* consists of *support* and *attack* relations on a set of arguments. Each argument is assigned a *basic weight* from the real interval $[0, 1]$.

Definition 1 (wBAG). A weighted bipolar argumentation graph (wBAG) is a quadruple $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$, where \mathcal{A} is a non-empty finite set of arguments, w is a function from \mathcal{A} to $[0, 1]$, $S \subseteq \mathcal{A} \times \mathcal{A}$ and $R \subseteq \mathcal{A} \times \mathcal{A}$.

Given two arguments $a, b \in \mathcal{A}$, $(a, b) \in S$ means a supports b and $(a, b) \in R$ means a attacks b . By $w(a)$ we denote the basic weight of a , which may represent various measures such as trustworthiness (da Costa Pereira, Tettamanzi, and Villata 2011), votes (Leite and Martins 2011), probability of beliefs (Hunter 2013), etc.

In this paper, arguments are evaluated through the *bilateral gradual semantics* which assigns each argument an *acceptability degree* and a *rejectability degree*.

Definition 2 (Bilateral gradual semantics). A bilateral gradual semantics is a function \mathcal{S} transforming $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ to a function $\text{Deg}_{\mathbf{G}}^{\mathcal{S}}$ defined from \mathcal{A} to $[0, 1] \times [0, 1]$. For any $a \in \mathcal{A}$, $\text{Deg}_{\mathbf{G}}^{\mathcal{S}}(a) = (\sigma_{\mathbf{G}}^+(a), \sigma_{\mathbf{G}}^-(a))$ where $\sigma_{\mathbf{G}}^+(a)$ and $\sigma_{\mathbf{G}}^-(a)$ represent the acceptability and rejectability degrees of a respectively.

When the context is clear, we simply write Deg (resp. σ^+ , σ^-) instead of $\text{Deg}_{\mathbf{G}}^{\mathcal{S}}$ (resp. $\sigma_{\mathbf{G}}^+$, $\sigma_{\mathbf{G}}^-$). Below, we present some notations used in the paper. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $a \in \mathcal{A}$. $\text{Sup}_{\mathbf{G}}(a)$ denotes the set of all supporters of a , i.e., $\text{Sup}_{\mathbf{G}}(a) = \{b \in \mathcal{A} \mid (b, a) \in S\}$. $\text{Att}_{\mathbf{G}}(a)$ denotes the set of all attackers of a , i.e., $\text{Att}_{\mathbf{G}}(a) = \{b \in \mathcal{A} \mid (b, a) \in R\}$. We say that a is *non-supported* if $\text{Sup}_{\mathbf{G}}(a) = \emptyset$ and *non-attacked* if $\text{Att}_{\mathbf{G}}(a) = \emptyset$. We may abbreviate $\text{Att}_{\mathbf{G}}(a)$ as $\text{Att}(a)$ and $\text{Sup}_{\mathbf{G}}(a)$ as $\text{Sup}(a)$. For $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', w', S', R' \rangle$ s.t. $\mathcal{A} \cap \mathcal{A}' = \emptyset$, we define $\mathbf{G} \oplus \mathbf{G}' = \langle \mathcal{A} \cup \mathcal{A}', w^*, S \cup S', R \cup R' \rangle$ where for any $a \in \mathcal{A}$ (resp. $a \in \mathcal{A}'$), $w^*(a) = w(a)$ (resp. $w^*(a) = w'(a)$).

Principles

Principles represent a set of desirable properties that semantics usually need to satisfy in practical applications, serving as a guideline for exploring semantics. Numerous studies have investigated the principles that consider only the acceptability degree (Amgoud and Ben-Naim 2018a; Baroni, Rago, and Toni 2018, 2019; Amgoud, Doder, and Vesic 2022; Bonzon et al. 2016). Principles for bilateral gradual semantics in unipolar attack graphs were proposed in (Wang and Shen 2024). Based on this bilateral philosophy, we study the principles that simultaneously take into account the acceptability and rejectability degrees in wBAG.

We provide a total of 23 principles, which integrate and adapt the well-established principles from unipolar support graphs (Amgoud and Ben-Naim 2016), unipolar attack graphs (Amgoud et al. 2017), and bipolar graphs (Amgoud and Ben-Naim 2018b). Some principles are basic ones, e.g., Anonymity, Resilience, Proportionality. There is also a set of symmetric principles that consider how supporters and attackers respectively influence the acceptability and rejectability degrees of arguments, e.g., A-Counting, R-Counting, A-Reinforcement, R-Reinforcement. We also incorporate three notable principles—*Quality Precedence*, *Cardinality Precedence* and *Compensation*—which serve as strategies for defining semantics in unipolar graphs, yet have never been studied in bipolar graphs.

Table 1 illustrates the sources of acceptability and rejectability degrees. Intuitively, the former is determined by supporters and basic weight, while the latter is solely determined by attackers. Such a bilateral and non-reciprocal setting originates from (Wang and Shen 2024) for unipolar graphs. Here we lift the idea for bipolar graphs, and this turns out to be a cornerstone for developing convergent semantics.

Degree	Source
acceptability	<ul style="list-style-type: none"> • basic weight • acceptability degree of supporters • rejectability degree of supporters
rejectability	<ul style="list-style-type: none"> • acceptability degree of attackers • rejectability degree of attackers

Table 1: Sources of acceptability and rejectability degrees

The notions of *Isomorphism* and *Path* will be used to establish principles.

Definition 3 (Isomorphism). Consider $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', w', S', R' \rangle$. An isomorphism from \mathbf{G} to \mathbf{G}' is a bijective function f from \mathcal{A} to \mathcal{A}' s.t. (i) $\forall a \in \mathcal{A}, w(a) = w'(f(a))$, (ii) $\forall a, b \in \mathcal{A}, (a, b) \in S \text{ iff } (f(a), f(b)) \in S'$, (iii) $\forall a, b \in \mathcal{A}, (a, b) \in R \text{ iff } (f(a), f(b)) \in R'$.

Definition 4 (Path). We say that there is a path from a_1 to a_n iff there is a sequence consisting of $\{a_1, \dots, a_n\}$ s.t. $(a_i, a_{i+1}) \in R \cup S$ for any $i \in \{1, \dots, n-1\}$.

Anonymity. The acceptability and rejectability degrees of an argument are independent of its identity. Formally:

(1) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', w', S', R' \rangle$, for any isomorphism f from \mathbf{G} to \mathbf{G}' , we have $\forall a \in \mathcal{A}, \sigma_{\mathbf{G}}^+(a) = \sigma_{\mathbf{G}'}^+(f(a)), \sigma_{\mathbf{G}}^-(a) = \sigma_{\mathbf{G}'}^-(f(a))$.

Independence. The acceptability and rejectability degrees of an argument should be independent of any argument that is not connected to it.

(2) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', w', S', R' \rangle$ s.t. $\mathcal{A} \cap \mathcal{A}' = \emptyset$, we have $\forall a \in \mathcal{A}, \sigma_{\mathbf{G}}^+(a) = \sigma_{\mathbf{G}' \oplus \mathbf{G}}^+(a), \sigma_{\mathbf{G}}^-(a) = \sigma_{\mathbf{G}' \oplus \mathbf{G}}^-(a)$.

Directionality. The acceptability and rejectability degrees of an argument a depend on argument b only if there is a path from b to a .

(3) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}, w, S', R' \rangle$ s.t. $S \subseteq S'$, $R \subseteq R'$, $R' \cup S' = R \cup S \cup \{(a, b)\}$, $\forall x \in \mathcal{A}$, if there is no path from b to x , then $\sigma_{\mathbf{G}}^+(x) = \sigma_{\mathbf{G}'}^+(x)$, $\sigma_{\mathbf{G}}^-(x) = \sigma_{\mathbf{G}'}^-(x)$.

Equivalence. The acceptability degree of an argument only depends on the acceptability and rejectability degrees of its supporters, as well as its basic weight; (ii) The rejectability degree only depends on the acceptability and rejectability degrees of its attackers.

(4) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$, $\forall a, b \in \mathcal{A}$,

- if $w(a) = w(b)$, and there exists a bijective function f from $\text{Sup}(a)$ to $\text{Sup}(b)$ s.t. $\forall x \in \text{Sup}(a)$, $\sigma^+(x) = \sigma^+(f(x))$, $\sigma^-(x) = \sigma^-(f(x))$, then $\sigma_{\mathbf{G}}^+(a) = \sigma_{\mathbf{G}}^+(b)$;
- if there exists a bijective function f' from $\text{Att}(a)$ to $\text{Att}(b)$ s.t. $\forall x \in \text{Att}(a)$, $\sigma^+(x) = \sigma^+(f'(x))$, $\sigma^-(x) = \sigma^-(f'(x))$, then $\sigma_{\mathbf{G}}^-(a) = \sigma_{\mathbf{G}}^-(b)$.

Resilience. If the basic weight of an argument lies in $(0, 1)$, then its acceptability degree also lies in $(0, 1)$.

(5) If $0 < w(a) < 1$, then $0 < \sigma^+(a) < 1$.

Proportionality. The higher the basic weight of an argument, the higher its acceptability degree.

(6) If (i) $w(a) < w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, and (iii) $\text{Sup}(a) = \text{Sup}(b)$, then $\sigma^+(a) < \sigma^+(b)$.

The following principles consider how supporters and attackers respectively determine the acceptability and rejectability degrees of arguments. We say an argument a is *worthless* if $\sigma^+(a) = 0$ and *alive* if $\sigma^+(a) > 0$.

A-Neutrality. A worthless supporter has no impact on the acceptability degree of the supported argument.

(7) If (i) $w(a) = w(b)$, (ii) $\text{Sup}(a) = \text{Sup}(b) \setminus \{x\}$ with $x \in \text{Sup}(b)$, (iii) $\sigma^+(x) = 0$, then $\sigma^+(a) = \sigma^+(b)$.

R-Neutrality. A worthless attacker has no impact on the rejectability degree of the attacked argument.

(8) If (i) $\text{Att}(a) = \text{Att}(b) \setminus \{x\}$ with $x \in \text{Att}(b)$, (ii) $\sigma^+(x) = 0$, then $\sigma^-(a) = \sigma^-(b)$.

A-Stability. The acceptability degree of a non-supported argument is equal to its basic weight.

(9) If $\text{Sup}(a) = \emptyset$, then $\sigma_{\mathbf{G}}^+(a) = w(a)$.

R-Stability. The rejectability degree of a non-attacked argument is 0.

(10) If $\text{Att}(a) = \emptyset$, then $\sigma_{\mathbf{G}}^-(a) = 0$.

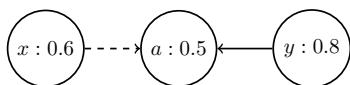
A-Strengthening. An alive supporter strengthens the acceptability degree of an argument to be greater than its basic weight.

(11) If $w(a) < 1$ and $\exists b \in \text{Sup}(a)$ s.t. $\sigma^+(b) > 0$, then $\sigma^+(a) > w(a)$.

R-Strengthening. An alive attacker strengthens the rejectability degree of an argument to be greater than 0.

(12) If $\exists b \in \text{Att}(a)$ s.t. $\sigma^-(b) > 0$, then $\sigma^-(a) > 0$.

Example 1. A wBAG where dashed arrows represent supports and solid arrows represent attacks. Arguments x, a, y are assigned with basic weights 0.6, 0.5, 0.8 respectively.



Assume a semantics satisfies A-Stability and R-Stability. Then $\sigma^+(x) = 0.6$, $\sigma^+(y) = 0.8$, and $\sigma^-(x) = \sigma^-(y) = 0$. By A-Strengthening and R-Strengthening, we have $\sigma^+(a) > 0.5$ and $\sigma^-(a) > 0$.

A-Strengthening Soundness. Alive supporters are the only source of obtaining the acceptability degree.

(13) If $w(a) < 1$ and $\sigma^+(a) > w(a)$, then $\exists b \in \text{Sup}(a)$ s.t. $\sigma^+(b) > 0$.

R-Strengthening Soundness. Alive attackers are the only source of obtaining the rejectability degree.

(14) If $\sigma^-(a) > 0$, then $\exists b \in \text{Att}(a)$ s.t. $\sigma^-(b) > 0$.

The following two principles claim that each alive argument has an impact on the arguments it attacks or supports.

A-Counting. Adding an alive supporter leads to an increase in the acceptability degree of the supported argument.

(15) If (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $\text{Sup}(a) = \text{Sup}(b) \setminus \{x\}$ with $x \in \text{Sup}(b)$, (iv) $\sigma^+(x) > 0$, then $\sigma^+(a) < \sigma^+(b)$.

R-Counting. Adding an alive attacker leads to an increase in the rejectability degree of the attacked argument.

(16) If (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $\text{Att}(a) = \text{Att}(b) \setminus \{x\}$ with $x \in \text{Att}(b)$, (iii) $\sigma^-(x) > 0$, then $\sigma^-(a) < \sigma^-(b)$.

Intuitively, the higher the acceptability degree, the stronger the impact of the argument.

A-Reinforcement. Increasing a supporter's acceptability degree leads to an increase in the acceptability degree of the supported argument.

(17) If (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $\text{Sup}(a) \setminus \{x\} = \text{Sup}(b) \setminus \{y\}$ with $x \in \text{Sup}(a)$ and $y \in \text{Sup}(b)$, (iv) $\sigma^+(x) < \sigma^+(y)$ and $\sigma^-(x) = \sigma^-(y)$, then $\sigma^+(a) < \sigma^+(b)$.

R-Reinforcement. Increasing an attacker's acceptability degree leads to an increase in the rejectability degree of the attacked argument.

(18) If (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $\text{Att}(a) \setminus \{x\} = \text{Att}(b) \setminus \{y\}$ with $x \in \text{Att}(a)$ and $y \in \text{Att}(b)$, (iii) $\sigma^-(x) < \sigma^-(y)$ and $\sigma^-(x) = \sigma^-(y)$, then $\sigma^-(a) < \sigma^-(b)$.

Naturally, the higher the rejectability degree, the weaker the impact of the argument.

A-Weakened Defense. Increasing a supporter's rejectability degree leads to a decrease in the acceptability degree of the supported argument.

(19) If (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $\text{Sup}(a) \setminus \{x\} = \text{Sup}(b) \setminus \{y\}$ with $x \in \text{Sup}(a)$ and $y \in \text{Sup}(b)$, (iv) $\sigma^-(x) < \sigma^-(y)$ and $\sigma^+(x) = \sigma^+(y) > 0$, then $\sigma^+(a) > \sigma^+(b)$.

R-Weakened Defense. Increasing an attacker's rejectability degree leads to a decrease in the rejectability degree of the attacked argument.

(20) If (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $\text{Att}(a) \setminus \{x\} = \text{Att}(b) \setminus \{y\}$ with $x \in \text{Att}(a)$ and $y \in \text{Att}(b)$, (iii) $\sigma^-(x) < \sigma^-(y)$ and $\sigma^+(x) = \sigma^+(y) > 0$, then $\sigma^-(a) > \sigma^-(b)$.

The last three notable principles correspond to three strategies to deal with the precedence of the *quality*, *quantity* or *compromise* of supporters and attackers.

Quality Precedence (QP) prioritizes the quality of supporters and attackers. It focuses on the *strongest* supporter

and attacker, which have both the highest acceptability and lowest rejectability degree. In words: (i) the stronger the strongest supporter of an argument, the higher its acceptability, and (ii) the stronger the strongest attacker of an argument, the higher its rejectability.

(21) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle, \forall a, b \in \mathcal{A}$,

- if (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $\exists y \in \text{Sup}(b)$ s.t. $\forall x \in \text{Sup}(a), \sigma^+(x) < \sigma^+(y)$ and $\sigma^-(x) > \sigma^-(y)$, then $\sigma^+(a) < \sigma^+(b)$;
- if (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $\exists y \in \text{Att}(b)$ s.t. $\forall x \in \text{Att}(a), \sigma^+(x) < \sigma^+(y)$ and $\sigma^-(x) > \sigma^-(y)$, then $\sigma^-(a) < \sigma^-(b)$.

Cardinality Precedence (CP) prioritizes the quantity of supporters and attackers. In words: (i) the greater the number of alive supporters of an argument, the higher its acceptability, and (ii) the greater the number of alive attackers of an argument, the higher its rejectability.

(22) For any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle, \forall a, b \in \mathcal{A}$,

- if (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $|\{x \in \text{Sup}(a) | \sigma^+(x) > 0\}| < |\{y \in \text{Sup}(b) | \sigma^+(y) > 0\}|$, then $\sigma^+(a) < \sigma^+(b)$;
- if (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $|\{x \in \text{Att}(a) | \sigma^+(x) > 0\}| < |\{y \in \text{Att}(b) | \sigma^+(y) > 0\}|$, then $\sigma^-(a) < \sigma^-(b)$.

Compensation is a kind of compromise that considers both the quality and quantity of attackers and supporters. It says that several weak supporters/attackers may compensate one strong supporter/attacker.

(23) There exists a wBAG $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ such that

- $\exists a, b \in \mathcal{A}$ s.t. (i) $w(a) = w(b)$, (ii) $\sigma^+(a) < 1$ or $\sigma^+(b) < 1$, (iii) $\exists y \in \text{Sup}(b)$ s.t. $\forall x \in \text{Sup}(a), \sigma^+(x) < \sigma^+(y)$ and $\sigma^-(x) > \sigma^-(y)$, (iv) $|\{x \in \text{Sup}(a) | \sigma^+(x) > 0\}| > |\{y \in \text{Sup}(b) | \sigma^+(y) > 0\}|$, (v) $\sigma^+(a) = \sigma^+(b)$;
- $\exists a, b \in \mathcal{A}$ s.t. (i) $\sigma^-(a) < 1$ or $\sigma^-(b) < 1$, (ii) $\exists y \in \text{Att}(b)$ s.t. $\forall x \in \text{Att}(a), \sigma^+(x) < \sigma^+(y)$ and $\sigma^-(x) > \sigma^-(y)$, (iii) $|\{x \in \text{Att}(a) | \sigma^+(x) > 0\}| > |\{y \in \text{Att}(b) | \sigma^+(y) > 0\}|$, (iv) $\sigma^-(a) = \sigma^-(b)$.

Formal Analysis of Principles

In this section, we provide a formal analysis of principles. We first present some links between principles.

Proposition 1. *The following properties hold:*

1. *A-Stability, R-Stability, R-Strengthening, QP and CP are incompatible.*
2. *Compensation is not compatible with QP or CP.*
3. *CP (resp. Compensation) is compatible with principles (1)-(20).*

Proposition 2. *Let \mathcal{S} be a semantics which satisfies Independence, Directionality and Equivalence. Then:*

1. *If \mathcal{S} satisfies A-Stability, A-Neutrality, then it also satisfies A-Strengthening Soundness.*
2. *If \mathcal{S} satisfies R-Stability, R-Neutrality, then it also satisfies R-Strengthening Soundness.*

3. *If \mathcal{S} satisfies A-Stability, A-Neutrality, A-Reinforcement, then it also satisfies A-Counting and A-Strengthening.*
4. *If \mathcal{S} satisfies A-Stability, R-Neutrality, R-Reinforcement, then it also satisfies R-Counting and R-Strengthening.*

Next, we study the behavior of semantics under some specified principles. To begin, we show that a set of arguments that are not supported or attacked by any other arguments keeps their acceptability and rejectability degrees unchanged in any graph, whenever the semantics satisfies Independence and Directionality.

Proposition 3. *If a semantics \mathcal{S} satisfies Independence and Directionality, then for any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', w', S', R' \rangle$ such that $\mathcal{A} \subseteq \mathcal{A}'$, $w(a) = w'(a)$ for all $a \in \mathcal{A}$, $S' \cap (\mathcal{A}' \times \mathcal{A}) = S$ and $R' \cap (\mathcal{A}' \times \mathcal{A}) = R$, it follows that $\sigma_{\mathbf{G}}^+(a) = \sigma_{\mathbf{G}'}^+(a)$ and $\sigma_{\mathbf{G}}^-(a) = \sigma_{\mathbf{G}'}^-(a)$ for all $a \in \mathcal{A}$.*

If an argument is only supported by worthless arguments, then its acceptability degree equals its basic weight, whenever the semantics satisfies Independence, Directionality, Equivalence, A-Stability and A-Neutrality.

Proposition 4. *Let \mathcal{S} be a semantics which satisfies Independence, Directionality, Equivalence, A-Stability and A-Neutrality. Then for any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle, \forall a \in \mathcal{A}$, if for any $x \in \text{Sup}(a), \sigma^+(x) = 0$, then $\sigma^+(a) = w(a)$.*

If an argument is only attacked by worthless arguments, then its rejectability degree is 0, whenever the semantics satisfies Independence, Directionality, Equivalence, R-Stability and R-Neutrality.

Proposition 5. *Let \mathcal{S} be a semantics which satisfies Independence, Directionality, Equivalence, R-Stability and R-Neutrality. Then for any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle, \forall a \in \mathcal{A}$, if for any $x \in \text{Att}(a), \sigma^+(x) = 0$, then $\sigma^-(a) = 0$.*

If a semantics satisfies Independence, Directionality, Equivalence, Proportionality, A-Neutrality, A-Stability and A-Strengthening, then the lower bound of the acceptability degree is equal to the basic weight.

Proposition 6. *Let \mathcal{S} be a semantics which satisfies Independence, Directionality, Equivalence, Proportionality, A-Neutrality, A-Stability and A-Strengthening. Then for any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle, \forall a \in \mathcal{A}, \sigma^+(a) \in [w(a), 1]$.*

Convergent Semantics and their Properties

In this section, we propose three convergent semantics for wBAG, each corresponding to one of QP, CP and Compensation. We develop semantics based on the well-studied *h-categorizer* function family (Besnard and Hunter 2001; Pu et al. 2014; Amgoud, Doder, and Vesic 2022). The resulting functions produce iterative sequences that always converge for any wBAG including cyclic ones.

Quality-Based Semantics

The *quality-based semantics* (QBS) prioritizes the quality of supporters and attackers, i.e., satisfies Quality Precedence.

Definition 5. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. For any $a \in \mathcal{A}$, the iterative sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ is defined by $F^i(a) = (f^i(a), g^i(a))$ where $f^i, g^i : \mathcal{A} \rightarrow [0, 1]$ are:

$$f^0(a) = w(a), g^0(a) = 0;$$

For any $i \geq 0$,

$$f^{i+1}(a) = w(a) + (1 - w(a)) \frac{\max_{b \in \text{Sup}(a)} h^i(b)}{1 + \max_{b \in \text{Sup}(a)} h^i(b)}$$

$$g^{i+1}(a) = \frac{\max_{b \in \text{Att}(a)} h^i(b)}{1 + \max_{b \in \text{Att}(a)} h^i(b)} \text{ with } h^i(b) = \frac{f^i(b)}{1 + g^i(b)}.$$

By convention, $\max_{b \in \text{Sup}(a)} h^i(b) = 0$ if $\text{Sup}(a) = \emptyset$ and $\max_{b \in \text{Att}(a)} h^i(b) = 0$ if $\text{Att}(a) = \emptyset$.

Theorem 1 (QBS Convergence). For any $a \in \mathcal{A}$, the sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ converges as i approaches infinity.

The quality-based semantics is defined through the limit of the above iterative sequence.

Definition 6 (QBS). The quality-based semantics is a function QBS transforming any $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ into a function $\text{Deg}_{\mathbf{G}}^{\text{QBS}}$ defined from \mathcal{A} to $[0, 1] \times [0, 1]$ s.t. $\forall a \in \mathcal{A}$, $\text{Deg}_{\mathbf{G}}^{\text{QBS}}(a) = (\sigma^+(a), \sigma^-(a))$ where $\sigma^+(a) = \lim_{i \rightarrow \infty} f^i(a)$, $\sigma^-(a) = \lim_{i \rightarrow \infty} g^i(a)$.

Theorem 2 states that the acceptability and rejectability degrees assigned by QBS can be nicely described by the following equations.

Theorem 2. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. Then for any $a \in \mathcal{A}$ under QBS, we have:

$$\sigma^+(a) = w(a) + (1 - w(a)) \frac{\max_{b \in \text{Sup}(a)} h(b)}{1 + \max_{b \in \text{Sup}(a)} h(b)}$$

$$\sigma^-(a) = \frac{\max_{b \in \text{Att}(a)} h(b)}{1 + \max_{b \in \text{Att}(a)} h(b)} \text{ with } h(b) = \frac{\sigma^+(b)}{1 + \sigma^-(b)}.$$

Theorem 3 states that QBS is the *unique* function that satisfies the above equations.

Theorem 3. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $D : \mathcal{A} \rightarrow [0, 1] \times [0, 1]$ be a function. If for any $a \in \mathcal{A}$, $D(a) = (D^+(a), D^-(a))$ satisfies

$$D^+(a) = w(a) + (1 - w(a)) \frac{\max_{b \in \text{Sup}(a)} h'(b)}{1 + \max_{b \in \text{Sup}(a)} h'(b)}$$

$$D^-(a) = \frac{\max_{b \in \text{Att}(a)} h'(b)}{1 + \max_{b \in \text{Att}(a)} h'(b)} \text{ with } h'(b) = \frac{D^+(b)}{1 + D^-(b)},$$

then $D = \text{Deg}_{\mathbf{G}}^{\text{QBS}}$.

Since QBS focuses on the strongest supporter and attacker, it violates A-Counting, R-Counting, A-Reinforcement, R-Reinforcement, A-Weakened Defense, and R-Weakened Defense. In fact, it satisfies the rest principles that are compatible with QP.

Theorem 4. QBS violates A-Counting, R-Counting, A-Reinforcement, R-Reinforcement, A-Weakened Defense, R-Weakened Defense, CP and Compensation. It satisfies all the remaining principles.

Cardinality-Based Semantics

The *cardinality-based semantics* (CBS) prioritizes the quantity of supporters and attackers, i.e., satisfies Cardinality Precedence. The iterative function below considers only the *founded* supporters/attackers whose previous values of the function f are greater than 0.

Definition 7. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. For any $a \in \mathcal{A}$, the iterative sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ is defined by $F^i(a) = (f^i(a), g^i(a))$ where $f^i, g^i : \mathcal{A} \rightarrow [0, 1]$ are:

$$f^0(a) = w(a), g^0(a) = 0;$$

For any $i \geq 0$,

$$f^{i+1}(a) = w(a) + (1 - w(a)) \frac{|\text{FS}^i(a)| + \frac{\sum_{b \in \text{FS}^i(a)} h^i(b)}{|\text{FS}^i(a)|}}{1 + |\text{FS}^i(a)| + \frac{\sum_{b \in \text{FS}^i(a)} h^i(b)}{|\text{FS}^i(a)|}}$$

$$g^{i+1}(a) = \frac{|\text{FA}^i(a)| + \frac{\sum_{b \in \text{FA}^i(a)} h^i(b)}{|\text{FA}^i(a)|}}{1 + |\text{FA}^i(a)| + \frac{\sum_{b \in \text{FA}^i(a)} h^i(b)}{|\text{FA}^i(a)|}},$$

in which $h^i(b) = \frac{f^i(b)}{1 + g^i(b)}$, $\text{FS}^i(a) = \{b \in \text{Sup}(a) | f^i(b) > 0\}$ and $\text{FA}^i(a) = \{b \in \text{Att}(a) | f^i(b) > 0\}$. We stipulate that $\frac{\sum_{b \in \text{FS}^i(a)} h^i(b)}{|\text{FS}^i(a)|} = 0$ if $\text{FS}^i(a) = \emptyset$ and $\frac{\sum_{b \in \text{FA}^i(a)} h^i(b)}{|\text{FA}^i(a)|} = 0$ if $\text{FA}^i(a) = \emptyset$.

Intuitively, $\text{FS}^i(a)$ (resp. $\text{FA}^i(a)$) denotes the set of founded supporters (resp. attackers) in the i -th iteration.

Theorem 5 (CBS Convergence). For any $a \in \mathcal{A}$, the sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ converges as i approaches infinity.

The cardinality-based semantics is defined through the limit of the above iterative sequence.

Definition 8 (CBS). The cardinality-based semantics is a function CBS transforming any wBAG $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ into a function $\text{Deg}_{\mathbf{G}}^{\text{CBS}}$ defined from \mathcal{A} to $[0, 1] \times [0, 1]$ s.t. $\forall a \in \mathcal{A}$, $\text{Deg}_{\mathbf{G}}^{\text{CBS}}(a) = (\sigma^+(a), \sigma^-(a))$ where $\sigma^+(a) = \lim_{i \rightarrow \infty} f^i(a)$, $\sigma^-(a) = \lim_{i \rightarrow \infty} g^i(a)$.

Theorem 6 states that the acceptability and rejectability degrees assigned by CBS can be nicely described by the following equations.

Theorem 6. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. Then for any $a \in \mathcal{A}$ under CBS, we have:

$$\sigma^+(a) = w(a) + (1 - w(a)) \frac{|\text{FS}(a)| + \frac{\sum_{b \in \text{FS}(a)} h(b)}{|\text{FS}(a)|}}{1 + |\text{FS}(a)| + \frac{\sum_{b \in \text{FS}(a)} h(b)}{|\text{FS}(a)|}}$$

$$\sigma^-(a) = \frac{|\text{FA}(a)| + \frac{\sum_{b \in \text{FA}(a)} h(b)}{|\text{FA}(a)|}}{1 + |\text{FA}(a)| + \frac{\sum_{b \in \text{FA}(a)} h(b)}{|\text{FA}(a)|}} \text{ with } h(b) = \frac{\sigma^+(b)}{1 + \sigma^-(b)},$$

$\text{FS}(a) = \{b \in \text{Sup}(a) \mid \sigma^+(b) > 0\}$ and $\text{FA}(a) = \{b \in \text{Att}(a) \mid \sigma^+(b) > 0\}$.

Similar to Theorem 3, Theorem 7 states that CBS is the unique function satisfying the above equations.

Theorem 7. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $D : \mathcal{A} \rightarrow [0, 1] \times [0, 1]$ be a function. If for any $a \in \mathcal{A}$, $D(a) = (D^+(a), D^-(a))$ satisfies the corresponding equations presented in Theorem 6. Then $D = \text{Deg}_{\mathbf{G}}^{\text{CBS}}$.

CBS satisfies all the principles that are compatible with CP, stated as below.

Theorem 8. CBS satisfies all the principles except QP and Compensation.

Hybrid-Based Semantics

The hybrid-based semantics (HBS) satisfies Compensation, taking both the quality and quantity into account.

Definition 9. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. For any $a \in \mathcal{A}$, the iterative sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ is defined by $F^i(a) = (f^i(a), g^i(a))$ where $f^i, g^i : \mathcal{A} \rightarrow [0, 1]$ are as follows:

$$f^0(a) = w(a), g^0(a) = 0;$$

For any $i \geq 0$,

$$f^{i+1}(a) = w(a) + (1 - w(a)) \frac{|\text{FS}^i(a)| + \sum_{b \in \text{FS}^i(a)} h^i(b)}{1 + |\text{FS}^i(a)| + \sum_{b \in \text{FS}^i(a)} h^i(b)}$$

$$g^{i+1}(a) = \frac{|\text{FA}^i(a)| + \sum_{b \in \text{FA}^i(a)} h^i(b)}{1 + |\text{FA}^i(a)| + \sum_{b \in \text{FA}^i(a)} h^i(b)},$$

in which $h^i(b) = \frac{f^i(b)}{1 + g^i(b)}$, $\text{FS}^i(a) = \{b \in \text{Sup}(a) \mid f^i(b) > 0\}$ and $\text{FA}^i(a) = \{b \in \text{Att}(a) \mid f^i(b) > 0\}$. We stipulate that $\sum_{b \in \text{FS}^i(a)} h^i(b) = 0$ if $\text{FS}^i(a) = \emptyset$ and $\sum_{b \in \text{FA}^i(a)} h^i(b) = 0$ if $\text{FA}^i(a) = \emptyset$.

Theorem 9 (HBS Convergence). For any $a \in \mathcal{A}$, the sequence $\{F^i(a)\}_{i \in \mathbb{N}}$ converges as i approaches infinity.

The hybrid-based semantics is defined through the limit of the above iterative sequence.

Definition 10 (HBS). The hybrid-based semantics is a function HBS transforming any wBAG $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ into a

function $\text{Deg}_{\mathbf{G}}^{\text{HBS}}$ defined from \mathcal{A} to $[0, 1] \times [0, 1]$ s.t. $\forall a \in \mathcal{A}$, $\text{Deg}_{\mathbf{G}}^{\text{HBS}}(a) = (\sigma^+(a), \sigma^-(a))$ where $\sigma^+(a) = \lim_{i \rightarrow \infty} f^i(a)$, $\sigma^-(a) = \lim_{i \rightarrow \infty} g^i(a)$.

Theorem 10 states that the acceptability and rejectability degrees assigned by HBS can be nicely described by the following equations.

Theorem 10. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$. Then for any $a \in \mathcal{A}$ under HBS, we have:

$$\sigma^+(a) = w(a) + (1 - w(a)) \frac{|\text{FS}(a)| + \sum_{b \in \text{FS}(a)} h(b)}{1 + |\text{FS}(a)| + \sum_{b \in \text{FS}(a)} h(b)}$$

$$\sigma^-(a) = \frac{|\text{FA}(a)| + \sum_{b \in \text{FA}(a)} h(b)}{1 + |\text{FA}(a)| + \sum_{b \in \text{FA}(a)} h(b)} \text{ with } h(b) = \frac{\sigma^+(b)}{1 + \sigma^-(b)},$$

$\text{FS}(a) = \{b \in \text{Sup}(a) \mid \sigma^+(b) > 0\}$ and $\text{FA}(a) = \{b \in \text{Att}(a) \mid \sigma^+(b) > 0\}$.

Similar to Theorem 3 and Theorem 7, Theorem 11 states that HBS is the unique function that satisfies the above equations.

Theorem 11. Let $\mathbf{G} = \langle \mathcal{A}, w, S, R \rangle$ and $D : \mathcal{A} \rightarrow [0, 1] \times [0, 1]$ be a function. If for any $a \in \mathcal{A}$, $D(a) = (D^+(a), D^-(a))$ satisfies the corresponding equations presented in Theorem 10, then $D = \text{Deg}_{\mathbf{G}}^{\text{HBS}}$.

HBS satisfies all the principles that are compatible with Compensation, stated as below.

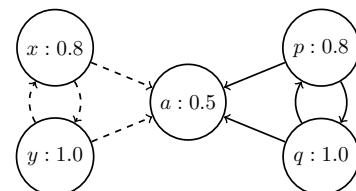
Theorem 12. HBS satisfies all the principles except QP and CP.

Comparisons for principles under QBS, CBS and HBS are summarized in Table 2.

	QBS	CBS	HBS
Principles (1)-(14)	•	•	•
A-Counting	–	•	•
R-Counting	–	•	•
A-Reinforcement	–	•	•
R-Reinforcement	–	•	•
A-Weakened Defense	–	•	•
R-Weakened Defense	–	•	•
Quality Precedence	•	–	–
Cardinality Precedence	–	•	–
Compensation	–	–	•

Table 2: Principles under QBS, CBS and HBS

Example 2. Consider the wBAG \mathbf{G} depicted below.



The strength of the argument a converges approximately as follows: $\text{Deg}_G^{\text{QBS}}(a) \approx (0.750, 0.424)$, $\text{Deg}_G^{\text{CBS}}(a) \approx (0.874, 0.720)$, and $\text{Deg}_G^{\text{HBS}}(a) \approx (0.899, 0.757)$.

Preliminary Experimental Analysis

We implemented our three novel semantics in Python and conducted preliminary experiments on the wBAG benchmark from (Potyka 2018). The benchmark¹ consists of 30 randomly generated (RG) datasets, each containing 100 graphs. The number of dataset arguments (i.e., graph size) ranges from 100 to 3000, increasing by 100.

The experiments were run on a Windows 11 x64 laptop with an Intel i7-11800H CPU (8 cores, 4.6 GHz) and 32 GB of memory. Our evaluation criteria include the *total running time* of a dataset and its *minimal/maximal iterations*. The algorithm stops when the difference of values between successive iterations < 0.0001 for each argument. All empirical results are summarized in Figure 1.

Generally speaking, the total running time appears to grow linearly in graph size and all semantics perform quite efficiently. The *average* time for computing a graph of 3000 arguments under QBS/HBS/CBS $\approx 0.11\text{s}/0.10\text{s}/0.09\text{s}$. Concerning the number of min/max iterations in the experiment, we observe that QBS usually takes more iterations to converge than HBS and CBS. Even so, QBS only uses 13 steps to converge in the worst case. Interestingly, the number of min/max iterations under HBS, CBS and QBS seems to be *constant* in graph size. A similar phenomenon is also mentioned in (Amgoud, Doder, and Vesic 2022). In addition, the min/max iterations of HBS and CBS appear to largely overlap. To sum up, the preliminary experiments suggest that our semantics perform rather efficiently and have potential for real-world applications.

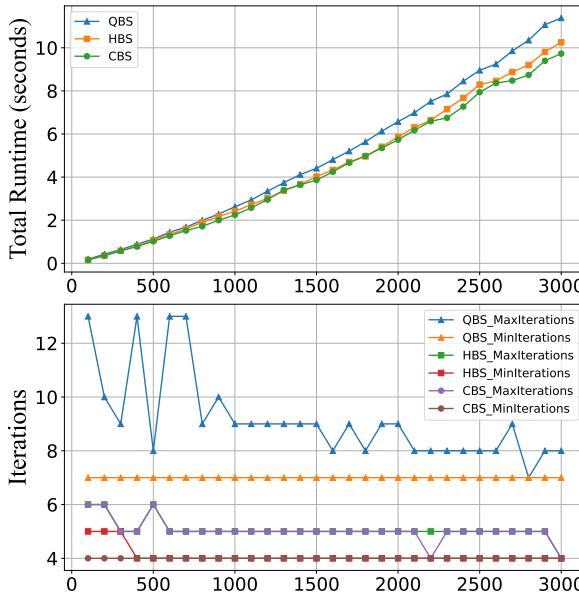


Figure 1: Total Runtime and Min/Max Iterations on RG

¹<https://www.researchgate.net/publication/326557254>

Discussion and Conclusion

The study of gradual semantics in wBAG has received extensive attention in the literature. Particularly, developing convergent semantics for cyclic wBAG is challenging. Most existing semantics are restricted to acyclic graphs, e.g., ESS-AAF (Evripidou and Toni 2014), QuAD (Baroni et al. 2015), DF-QuAD (Rago et al. 2016), Euler (Amgoud and Ben-Naim 2018b), MLP (Potyka 2021), and NPE (Doder, Vesic, and Croitoru 2021). The only exception is the Max Euler-based semantics (Mossakowski and Neuhaus 2018), which simply focuses on the strongest supporter and attacker while ignoring other arguments. This violates the so-called *open-mindedness* (Potyka 2019a), i.e., the strength of arguments cannot be far from their basic weights. In contrast, under our CBS and HBS, adding sufficient supporters (resp. attackers) leads to $\sigma^+ \rightarrow 1$ (resp. $\sigma^- \rightarrow 1$), which illustrates a form of open-mindedness.

The investigation of convergence behavior is of significant interest. Some special cyclic cases have been proven to converge under various semantics, typically by imposing constraints on the indegrees or basic weights of arguments (Mossakowski and Neuhaus 2018; Potyka 2021, 2019a; Potyka and Booth 2024a). Moreover, *continuization* has emerged as a technique to improve the convergence behavior of semantics (Potyka 2018, 2019a; Potyka and Booth 2024b). This approach associates update functions with a system of differential functions, and empirical studies have shown that continuized semantics are effective in resolving many divergence cases. However, their theoretical convergence behavior remains largely unknown.

The existing methodologies to aggregate the strength of supporters and attackers make it inherently difficult to obtain convergent semantics, as their aggregation operators do not always behave well during iterations. Particularly, in computing cyclic graphs with intricate interactions, argument values may vary drastically, resulting in divergent behavior (Amgoud and Ben-Naim 2018b; Potyka 2019a; Mossakowski and Neuhaus 2018; Potyka and Booth 2024a).

The paper elegantly addresses the challenge by adopting the bilateral gradual semantics, which offers a novel perspective to separately evaluate supporters and attackers through the acceptability and rejectability degrees. We first presented a set of principles for well-behaved wBAG semantics, and studied their relationships and properties. Then we proposed three convergent semantics that conform to the above principles. These semantics are defined through the limits of the iterative functions that always converge for any wBAG including cyclic ones, and perform efficiently on the benchmark from the literature.

Future work can be considered in several directions. We can investigate the computational complexity of convergent functions and conduct more experiments to check scalability. Additional desirable wBAG principles, such as monotonicity (Baroni, Rago, and Toni 2019), duality (Potyka 2020) can be discussed. It would also be interesting to aggregate σ^+ and σ^- to an overall strength for specific applications. Finally, applying wBAG to practical scenarios involving cycles of supports and attacks—such as social media debates and explainable AI—is a desirable endeavour.

Acknowledgements

This work was supported by the National Social Science Fund of China (25BZX071).

References

Amgoud, L.; and Ben-Naim, J. 2016. Evaluation of arguments from support relations: Axioms and semantics. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI*, 900–906.

Amgoud, L.; and Ben-Naim, J. 2018a. Evaluation of arguments in weighted bipolar graphs. *International Journal of Approximate Reasoning*, 99: 39–55.

Amgoud, L.; and Ben-Naim, J. 2018b. Weighted Bipolar Argumentation Graphs: Axioms and Semantics. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI*, 5194–5198.

Amgoud, L.; Ben-Naim, J.; Doder, D.; and Vesic, S. 2017. Acceptability Semantics for Weighted Argumentation Frameworks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI*, 56–62.

Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems*, 23(10): 1062–1093.

Amgoud, L.; Doder, D.; and Vesic, S. 2022. Evaluation of argument strength in attack graphs: Foundations and semantics. *Artificial Intelligence*, 302: 103607.

Amgoud, L.; and Prade, H. 2009. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173(3–4): 413–436.

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G.; Thimm, M.; and Villata, S. 2017. Towards artificial argumentation. *AI Magazine*, 38(3): 25–36.

Baroni, P.; Rago, A.; and Toni, F. 2018. How many properties do we need for gradual argumentation? In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI*, 126–139.

Baroni, P.; Rago, A.; and Toni, F. 2019. From fine-grained properties to broad principles for gradual argumentation: a principled spectrum. *International Journal of Approximate Reasoning*, 105: 252–286.

Baroni, P.; Romano, M.; Toni, F.; Aurisicchio, M.; and Bertranza, G. 2015. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation*, 6(1): 24–49.

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

Bonzon, E.; Delobelle, J.; Konieczny, S.; and Maudet, N. 2016. A comparative study of ranking-based semantics for abstract argumentation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI*, 914–920.

Cacioppo, J.; Gardner, W.; and Berntson, G. 1997. Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review*, 1(1): 3–25.

Cayrol, C.; and Lagasquie-Schiex, M. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU*, 378–389.

da Costa Pereira, C.; Tettamanzi, A.; and Villata, S. 2011. Changing one’s mind: Erase or rewind? Possibilistic belief revision with fuzzy argumentation based on trust. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI*, 164–171.

Doder, D.; Vesic, S.; and Croitoru, M. 2021. Ranking semantics for argumentation systems with necessities. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence, IJCAI*, 1912–1918.

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

Evripidou, V.; and Toni, F. 2014. Quaestio-it.com: a social intelligent debating platform. *Journal of Decision Systems*, 23(3): 333–349.

Gabbay, D.; and Rodrigues, O. 2015. Equilibrium states in numerical argumentation networks. *Logica Universalis*, 9(4): 411–473.

Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning*, 54(1): 47–81.

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

Mossakowski, T.; and Neuhaus, F. 2018. Modular semantics and characteristics for bipolar weighted argumentation graphs. *arXiv preprint arXiv:1807.06685*.

Polberg, S.; and Hunter, A. 2018. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning*, 93: 487–543.

Potyka, N. 2018. Continuous dynamical systems for weighted bipolar argumentation. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning, KR*, 148–157.

Potyka, N. 2019a. Extending Modular Semantics for Bipolar Weighted Argumentation. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 1722–1730.

Potyka, N. 2019b. Open-mindedness of gradual argumentation semantics. In *Proceedings of the 13th International Conference on Scalable Uncertainty Management, SUM*, 236–249.

Potyka, N. 2020. Bipolar abstract argumentation with dual attacks and supports. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR*, 677–686.

Potyka, N. 2021. Interpreting neural networks as quantitative argumentation frameworks. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence, AAAI*, 6463–6470.

Potyka, N.; and Booth, R. 2024a. Balancing open-mindedness and conservativeness in quantitative bipolar argumentation (and how to prove semantical from functional properties). In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning, KR*, 597–607.

Potyka, N.; and Booth, R. 2024b. An empirical study of quantitative bipolar argumentation frameworks for truth discovery. In *Proceedings of the 10th International Conference on Computational Models of Argument, COMMA*, 205–216. IOS Press.

Prakken, H. 2024. An abstract and structured account of dialectical argument strength. *Artificial Intelligence*, 335: 104193.

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

Rago, A.; Cocarascu, O.; Oksanen, J.; and Toni, F. 2025. Argumentative review aggregation and dialogical explanations. *Artificial Intelligence*, 340: 104291.

Rago, A.; and Toni, F. 2017. Quantitative argumentation debates with votes for opinion polling. In *Proceedings of the 20th International Conference on Principles and Practice of Multi-Agent Systems, PRIMA*, 369–385.

Rago, A.; Toni, F.; Aurisicchio, M.; and Baroni, P. 2016. Discontinuity-Free Decision Support with Quantitative Argumentation Debates. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning, KR*, 63–73.

Rago, A.; Vasileiou, S.; Toni, F.; Son, T.; and Yeoh, W. 2024. A Methodology for Gradual Semantics for Structured Argumentation under Incomplete Information. *arXiv preprint arXiv:2410.22209*.

Wang, Z.; and Shen, Y. 2024. Bilateral Gradual Semantics for Weighted Argumentation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence, AAAI*, 10732–10739.

Yin, X.; Potyka, N.; and Toni, F. 2024. CE-QArg: Counterfactual Explanations for Quantitative Bipolar Argumentation Frameworks. In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning, KR*, 697–707.