Computational Models of Argument
H. Prakken et al. (Eds.)

doi:10.3233/FAIA200492

PageRank as an Argumentation Semantics

Emanuele ALBINI [a,b,1], Pietro BARONI [a], Antonio RAGO [b] and Francesca TONI [b]
[a] *Dip.to di Ingegneria dell'Informazione, Università degli Studi di Brescia, Italy*
[b] *Dept. of Computing, Imperial College London, UK*

Abstract. This paper provides an initial exploration on the relationships between *PageRank* and gradual argumentation semantics. After showing that PageRank, directly interpreted as an argumentation semantics for support frameworks, fails to satisfy some generally desirable properties, we propose a novel approach to reconstruct PageRank as gradual semantics of a suitably defined bipolar argumentation framework, while satisfying these desirable properties. The theoretical advantages of the approach are complemented by an illustration of its potential application to support the generation of better explanations of PageRank scores for end users.

Keywords. PageRank, Gradual Argumentation Semantics, Bipolar Argumentation Frameworks

1. Introduction

In the context of search engines, a user wants to find the (web) pages that are the most relevant to a search query, potentially among millions of them. The web has an essential feature: each piece of information (page) may link to other pieces of information (through hyperlinks), and therefore the web organization can be regarded as a directed graph, where pages are nodes and links are the edges. This is the idea that in 1999 inspired the revolutionary PageRank (PR) algorithm [1]: a method for computing a ranking score for every page based on the graph structure of the web. Given its conceptual simplicity and general formalization for any kind of directed graph, PR has been applied to many other domains where entities can be evaluated on the basis of their connections to other entities, including citation networks [2], recommendation systems [3], chemistry [4], biology [5] and neuroscience [6], and has been studied from several perspectives including an axiomatic characterization from a social choice theory perspective [7].

As well-known, graph-based representations are also pervasive in the field of computational argumentation. In particular Dung's abstract argumentation frameworks [8] are essentially directed graphs whose nodes are arguments and edges represent attacks. Dung's seminal proposal has been subsequently extended in several directions, e.g. bipolar argumentation frameworks [9] encompass also a notion of support, while in quantitative bipolar argumentation frameworks [10] a base score is assigned to each argument. In this context, the argument graph structure is the basis of the assessment of argument acceptability according to some *argumentation semantics* [11]: in Dung's traditional approach the evaluation is qualitative, while in further developments numerical argument assessments based on *gradual semantics* have been investigated [12,10].

[1] Corresponding Author: emanuele.albini19@imperial.ac.uk

Given the similarity between PR and gradual argumentation semantics as formal tools producing a numerical assessment of connected entities in a graph, it appears that drawing bridges between the two areas and exploring possible cross-fertilization opportunities represents an interesting research direction. This paper provides some initial contribution in this respect by first exploring the use of PR as a gradual semantics for *support argumentation frameworks* [13], then evidencing some limitations of this simplistic correspondence and proposing a novel approach to reconstruct PR as a semantics in suitably constructed *quantitative bipolar argumentation frameworks* (QBAFs) in which pages will be interpreted as arguments ignoring their content. Besides featuring better theoretical properties, this approach has the significant advantage of supporting more effective explanations of PR outcomes to users.

In a broader perspective this paper contribution is two-fold. On one hand we define a new gradual semantics for QBAFs based on PageRank. On the other hand, we support the idea of using argumentation frameworks, not only to model dialectical debates, but also to describe the mechanism of algorithms in order to present them in a dialectical form, with the aim of either generating explanations or enabling other practical applications.

The paper is organized as follows. In Section 2 we recall some background concepts on PR. In Section 3 we detail how PR can be directly interpreted as a gradual semantics in support argumentation frameworks, showing however that, as such, it does not satisfy some desirable properties in this context. In Section 4 we reconstruct PR as a gradual semantics of suitable QBAFs, achieving in this way the satisfaction of the above mentioned desirable properties. In Section 5 we discuss the advantages of the proposed approach, with particular reference to the explanation of PR outcomes. We conclude the paper and outline lines of future work in Section 6.

2. PageRank Background

We firstly recall the PR definition from the original paper [1], using a different but equivalent notation when necessary for our purposes.

We assume a set of pages/nodes $\mathcal{P} = \{u_1, u_2, ..., u_N\}$ and a set of links between the pages $\mathcal{L} \subseteq \mathcal{P} \times \mathcal{P}$, where $(u, v) \in \mathcal{L}$ indicates that there is a link from page u to page v. We say $N = |\mathcal{P}| > 0$ is the total number of pages, $O_u = \{v \in \mathcal{P}:(u, v) \in \mathcal{L}\}$ is the set of pages u points to and $I_u = \{v \in \mathcal{P}:(v, u) \in \mathcal{L}\}$ is the set of pages that point to u. We assume that $\forall u \in \mathcal{P}, \nexists(u, u) \in \mathcal{L}$, i.e. self-loops are ignored to prevent the manipulation of PR. We also assume that $\forall u \in \mathcal{P}, |O_u| > 0$, i.e. there are no *dangling* pages, that is, no pages without outgoing links (in practice, if such a page is found it is treated as having links towards all other pages as in [14]).

A *random surfer model* is used, which is based on the assumption that a user can either reach a page from a link in another page with probability $d \in]0; 1[$, referred to as *damping factor*, or land on a page directly with probability $1 - d$. Unless otherwise specified, we assume the value suggested in [1] of $d = 0.85$ and a uniform probability of directly landing on a page (i.e. we focus on *non-personalized PR*). In Section 6 we discuss how in future works these assumptions could be changed.

Definition 1. *[1] The* PageRank *of a set of pages is an assignment,* $R : \mathcal{P} \to]0, 1]$*, to the pages which satisfies:*

$$R(u) = (1-d) \cdot \frac{1}{N} + d \cdot \sum_{v \in I_u} \frac{R(v)}{|O_v|} \quad \forall u \in \mathcal{P}$$

Note that R is the solution of a system of linear equations derived from Def. 1 (we refer to R as both the assignment and the vector resulting from it). Notice also that, as described in [14], R is unique and $||R||_1 = 1$, i.e. the L_1 norm of R is 1.

The aim of PR is to give to every page a score that describes how relevant it is: the higher the score, the more important the page. Thus, these scores are based on their relevance, which is intended to approximate the amount of users visiting the page. The latter is calculated through a mathematical model aiming at probabilistically estimating the number of users. The assumption here is therefore that the higher the number of links to (from) a page, the more it (the less each page linked by it, resp.) will be visited and hence the higher (lower, resp.) its PR score should be.

3. PageRank as a Gradual Semantics

In this section we show how PR may be interpreted directly as a gradual argumentation semantics and examine its ability to satisfy some desirable properties. First, we recall some necessary formal notions from [10].

Definition 2. *[10] A* Quantitative Bipolar Argumentation Framework *(QBAF) is a 4-tuple $\langle \mathcal{X}, \mathcal{R}^-, \mathcal{R}^+, \tau \rangle$, comprising:*

- *a finite set of arguments $\mathcal{X}$*
- *a binary attack relation between arguments $\mathcal{R}^- \subseteq \mathcal{X} \times \mathcal{X}$*
- *a binary support relation between arguments $\mathcal{R}^+ \subseteq \mathcal{X} \times \mathcal{X}$*
- *a total function $\tau : \mathcal{X} \to \mathbb{I}$, with $\tau(\alpha)$ the* base score *of α*

where $\mathbb{I}$ is a set equipped with a preorder $\leq$ where, as usual, $a<b$ denotes $a \leq b$ and $b \nleq a$. Given a QBAF, a total function $\sigma : \mathcal{X} \to \mathbb{I}$, called a gradual semantics*, may be used to assign a* strength *to each argument. We define an sQBAF as a QBAF such that $\mathcal{R}^- = \emptyset$. Finally, we let $\mathcal{R}^-(\alpha) = \{\beta \in \mathcal{X} : (\beta, \alpha) \in \mathcal{R}^-\}$ and $\mathcal{R}^+(\alpha) = \{\beta \in \mathcal{X} : (\beta, \alpha) \in \mathcal{R}^+\}$, and similarly $\mathfrak{R}^-(\alpha) = \{\beta \in \mathcal{X} : (\alpha, \beta) \in \mathcal{R}^-\}$ and $\mathfrak{R}^+(\alpha) = \{\beta \in \mathcal{X} : (\alpha, \beta) \in \mathcal{R}^+\}$.*

A *web graph* $\langle \mathcal{P}, \mathcal{L} \rangle$ can be interpreted as an sQBAF where the pages (nodes) are arguments and the links between them (edges) are supports, as follows.

Definition 3. *Given a set of pages $\mathcal{P}$ and a set of links $\mathcal{L}$, a* PageRank Argumentation Framework *(PRAF) is an sQBAF defined as $PR = \langle \mathcal{X}, \emptyset, \mathcal{R}^+, \tau \rangle$, where:*

- *$\mathcal{X} = \mathcal{P}$ is the set of arguments corresponding to the set of pages*
- *$\mathcal{R}^+ = \mathcal{L}$ is the set of supports corresponding to the set of links between pages*
- *$\tau : \mathcal{X} \mapsto \mathbb{I} = [\frac{1-d}{|\mathcal{X}|}, 1]$ is the base score, defined as a constant function:*

$$\tau(\alpha) = \frac{1-d}{|\mathcal{X}|} \quad \forall \alpha \in \mathcal{X}$$

Given Def. 1 and the notes on loops and dangling nodes in Section 2, Remark 1 can be trivially derived.

Remark 1. *Given a PRAF it always holds that:*
- *each argument has at least one outgoing link:* $|\mathfrak{R}^{+}(\alpha)|>0,\ \forall\alpha\in\mathcal{X}$
- *there are no self-supports:* $\nexists(\alpha,\alpha)\in\mathcal{R}^{+},\ \forall\alpha\in\mathcal{X}$.

We then interpret PR as a gradual semantics for sQBAF.

Definition 4. *The PageRank semantics is a gradual semantics* $\sigma:\mathcal{X}\mapsto\mathbb{I}$ *such that:*

$$\sigma(\alpha)=\tau(\alpha)+d\cdot\sum_{\beta\in\mathcal{R}^{+}(\alpha)}\frac{\sigma(\beta)}{|\mathfrak{R}^{+}(\beta)|}\quad\forall\alpha\in\mathcal{X}$$

The following lemma is directly derived from Def. 4.

Lemma 1. *The codomain of* σ *is* $\mathbb{I}=[\frac{1-d}{|\mathcal{X}|},1]$ *where* $\bot>0$.

In order to formally assess PR as an argumentation semantics, we now review some desirable properties for argument strength, called *group properties* (GPs) in [10], as they imply a group of other properties. Some preliminary definitions need to be recalled first. Given a QBAF $\langle\mathcal{X},\mathcal{R}^{-},\mathcal{R}^{+},\tau\rangle$ and a gradual semantics σ, for any $A\subseteq\mathcal{X}$, we refer to the multiset $\{\sigma(\beta){:}\beta\in A\}$ as A_σ. Given $A,B\subseteq\mathcal{X}$, A is *strength equivalent to* B, denoted $A\overset{\sigma}{=}B$, iff $A_\sigma=B_\sigma$; A is *at least as strong as* B, denoted $A\overset{\sigma}{\geq}B$, iff there exists an injective mapping f from B to A such that $\forall\alpha\in B$, $\sigma(f(\alpha))\geq\sigma(\alpha)$; and A is *stronger than* B, denoted $A\overset{\sigma}{>}B$, iff $A\overset{\sigma}{\geq}B$ and $B\overset{\sigma}{\ngeq}A$.

GPs are then defined as follows (some being reformulated in more general or more specific ways wrt [10], where useful for our present purposes):

GP1. If $\mathcal{R}^{-}(\alpha)=\emptyset$ and $\mathcal{R}^{+}(\alpha)=\emptyset$ then $\sigma(\alpha)=\tau(\alpha)$.
GP2. If $\mathcal{R}^{-}(\alpha)\neq\emptyset$ and $\mathcal{R}^{+}(\alpha)=\emptyset$ then $\sigma(\alpha)<\tau(\alpha)$.
GP3. If $\mathcal{R}^{-}(\alpha)=\emptyset$ and $\mathcal{R}^{+}(\alpha)\neq\emptyset$ then $\sigma(\alpha)>\tau(\alpha)$.
GP4. If $\sigma(\alpha)<\tau(\alpha)$ then $\mathcal{R}^{-}(\alpha)\neq\emptyset$.
GP5. If $\sigma(\alpha)>\tau(\alpha)$ then $\mathcal{R}^{+}(\alpha)\neq\emptyset$.
GP6. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{=}\mathcal{R}^{-}(\beta)$, $\mathcal{R}^{+}(\alpha)\overset{\sigma}{=}\mathcal{R}^{+}(\beta)$ and $\tau(\alpha)=\tau(\beta)$ then $\sigma(\alpha)=\sigma(\beta)$.
GP7. If $\mathcal{R}^{-}_{\sigma}(\alpha)\subsetneq\mathcal{R}^{-}_{\sigma}(\beta)$, $\mathcal{R}^{+}(\alpha)\overset{\sigma}{=}\mathcal{R}^{+}(\beta)$ and $\tau(\alpha)=\tau(\beta)$ then $\sigma(\beta)<\sigma(\alpha)$.
GP8. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{=}\mathcal{R}^{-}(\beta)$, $\mathcal{R}^{+}_{\sigma}(\alpha)\subsetneq\mathcal{R}^{+}_{\sigma}(\beta)$ and $\tau(\alpha)=\tau(\beta)$ then $\sigma(\alpha)<\sigma(\beta)$.
GP9. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{=}\mathcal{R}^{-}(\beta)$, $\mathcal{R}^{+}(\alpha)\overset{\sigma}{=}\mathcal{R}^{+}(\beta)$ and $\tau(\alpha)<\tau(\beta)$ then $\sigma(\alpha)<\sigma(\beta)$.
GP10. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{<}\mathcal{R}^{-}(\beta)$, $\mathcal{R}^{+}(\alpha)\overset{\sigma}{=}\mathcal{R}^{+}(\beta)$ and $\tau(\alpha)=\tau(\beta)$ then $\sigma(\beta)<\sigma(\alpha)$.
GP11. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{=}\mathcal{R}^{-}(\beta)$, $\mathcal{R}^{+}(\alpha)\overset{\sigma}{>}\mathcal{R}^{+}(\beta)$ and $\tau(\alpha)=\tau(\beta)$ then $\sigma(\beta)<\sigma(\alpha)$.

In [10], two general principles (and their strict counterparts) were also identified as a more synthetic way of describing the desirable properties of a gradual semantics.

The intuition for the first principle is that a difference in an argument's strength and base score must correspond to an imbalance in its attackers' and supporters' strengths.

Principle 1. *[10] A gradual semantics* σ *is* balanced *iff for any* $\alpha\in\mathcal{X}$*:*
1. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{=}\mathcal{R}^{+}(\alpha)$ *then* $\sigma(\alpha)=\tau(\alpha)$.
2. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{>}\mathcal{R}^{+}(\alpha)$ *then* $\sigma(\alpha)<\tau(\alpha)$.
3. If $\mathcal{R}^{-}(\alpha)\overset{\sigma}{<}\mathcal{R}^{+}(\alpha)$ *then* $\sigma(\alpha)>\tau(\alpha)$.

A gradual semantics σ is strictly balanced *iff σ is balanced and for any $\alpha \in \mathcal{X}$:*

4. If $\sigma(\alpha) < \tau(\alpha)$ then $\mathcal{R}^-(\alpha) \overset{\sigma}{>} \mathcal{R}^+(\alpha)$.

5. If $\sigma(\alpha) > \tau(\alpha)$ then $\mathcal{R}^-(\alpha) \overset{\sigma}{<} \mathcal{R}^+(\alpha)$.

In [10] it is shown that if σ is balanced then it satisfies GP1 to GP3 and if it is strictly balanced then it satisfies GP1 to GP5.

The second principle requires that the strength of an argument depends monotonically on its base score and on the strengths of its attackers and supporters. To introduce this principle formally, we first recall the notion of shaping triple of an argument, where for any $\alpha \in \mathcal{X}$, the *shaping triple* of α is $(\tau(\alpha), \mathcal{R}^+(\alpha), \mathcal{R}^-(\alpha))$, denoted $\mathcal{ST}(\alpha)$. Given $\alpha, \beta \in \mathcal{X}$, $\mathcal{ST}(\beta)$ is said to be: *as boosting as $\mathcal{ST}(\alpha)$*, denoted as $\mathcal{ST}(\alpha) \simeq \mathcal{ST}(\beta)$, iff $\tau(\alpha) = \tau(\beta)$, $\mathcal{R}^+(\alpha) \overset{\sigma}{=} \mathcal{R}^+(\beta)$, and $\mathcal{R}^-(\beta) \overset{\sigma}{=} \mathcal{R}^-(\alpha)$; *at least as boosting as $\mathcal{ST}(\alpha)$*, denoted as $\mathcal{ST}(\alpha) \preceq \mathcal{ST}(\beta)$, iff $\tau(\alpha) \leq \tau(\beta)$, $\mathcal{R}^+(\alpha) \overset{\sigma}{\leq} \mathcal{R}^+(\beta)$, and $\mathcal{R}^-(\beta) \overset{\sigma}{\leq} \mathcal{R}^-(\alpha)$; or *strictly more boosting than $\mathcal{ST}(\alpha)$*, denoted as $\mathcal{ST}(\alpha) \prec \mathcal{ST}(\beta)$, iff $\mathcal{ST}(\alpha) \preceq \mathcal{ST}(\beta)$ and $\mathcal{ST}(\beta) \npreceq \mathcal{ST}(\alpha)$.

Principle 2. *[10] A gradual semantics σ is* monotonic *iff:*

1. for any $\alpha, \beta \in \mathcal{X}$, if $\mathcal{ST}(\alpha) \simeq \mathcal{ST}(\beta)$ then $\sigma(\alpha) = \sigma(\beta)$;

2. if $\mathcal{ST}(\alpha) \preceq \mathcal{ST}(\beta)$ then $\sigma(\alpha) \leq \sigma(\beta)$.

A gradual semantics σ is strictly monotonic iff σ is monotonic and:

3. for any $\alpha, \beta \in \mathcal{X}$, if $\mathcal{ST}(\alpha) \prec \mathcal{ST}(\beta)$ then $\sigma(\alpha) < \sigma(\beta)$.

In [10] it is shown that if σ is (strictly) monotonic then it satisfies GP6 to GP11.

We will now show that the PR semantics σ satisfies some, but not all, of the desirable properties for gradual semantics. We will consider whether or not the properties are satisfied by the semantics σ when applied to a generic QBAF, in Proposition 1 and 2, or when applied to a PRAF (denoted as $\langle PR, \sigma \rangle$), in Proposition 3 and 4 (see Table 1 for a compact summary). Note that in the first case, if attacks are present in the QBAF, they are simply ignored by the definition of the semantics, and some of the properties may not hold for this mere reason. Proofs for the *satisfied* GPs and principles have been omitted for lack of space, as they are not essential for this paper.

Proposition 1. *σ satisfies GP1, GP3, GP4, GP5 but not GP2, and thus is not balanced.*

Proof. GP2 does not hold as when $\mathcal{R}^+(\alpha) = \emptyset$, $\sigma(\alpha) = \tau(\alpha)$ independently of $\mathcal{R}^-(\alpha)$, which is ignored in the definition of σ. □

Proposition 2. *σ satisfies GP8 and GP9 but not GP6, GP7, GP10 and GP11, and thus is not monotonic.*

Proof. GP6: in the framework in Figure 1, we have $\mathcal{R}^+(\beta) \overset{\sigma}{=} \mathcal{R}^+(\delta)$ but $\sigma(\beta) \neq \sigma(\delta)$. GP7 and GP10 cannot hold as attackers do not affect σ. GP11: in the framework in Figure 1, we have $\mathcal{R}^+(\zeta) \overset{\sigma}{>} \mathcal{R}^+(\eta)$ but $\sigma(\zeta) < \sigma(\eta)$. □

Proposition 3. *$\langle PR, \sigma \rangle$ is strictly balanced and thus satisfies GP1 to GP5.*

Proposition 4. *$\langle PR, \sigma \rangle$ satisfies GP7 to GP10 but not GP6 or GP11 (provable by the counter-examples in Proposition 2), and thus is not monotonic.*

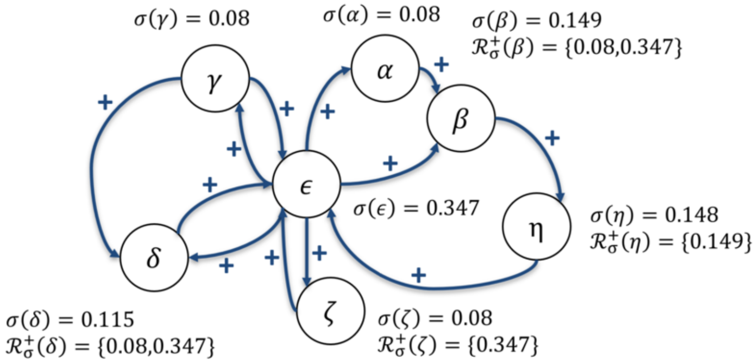

Figure 1. Counter-example to GP6 and GP11 for the PR semantics σ in Proposition 2.

We have thus shown that directly interpreting PR as a gradual semantics for an sQBAF does not give rise to a satisfactory outcome in terms of formal properties. Indeed, while using PR as a semantics is somehow straightforward, it does not appear fully appropriate from a modeling perspective, as it does not provide a suitable argumentative counterpart to some key aspects of PR. In particular, note that, as a consequence of the PR definition, the strength of each node depends not only on the strengths of its supporters but also on the cardinality of their outgoing supports. This has quite counter-intuitive effects from an argumentation perspective. For example, consider the situation where two nodes have the same strength $\sigma(\alpha) = \sigma(\beta)$, but α has one outgoing support, while β has ten: the latter's support to each of its children is actually ten times 'less powerful' (i.e. it transfers $1/10$ of the strength) than the former's. It follows that a node γ supported by α only and a node δ supported by β only would have different strengths even if their supporters appear to be equivalent (formally the shaping triples of γ and δ are the same). This is the main reason for the lack of many desirable properties and calls for an alternative approach, which we introduce next.

4. PageRank as a Gradual Semantics in a Meta-Argumentation Framework

In this section, we introduce an alternative approach to capture PageRank as an argumentation semantics. To this purpose we transform the sQBAF corresponding to a set of linked pages into a QBAF including additional meta-arguments and attacks between them. The underlying intuition is that each additional meta-argument can be understood as a vehicle of support from one page to another and that supports from the same page are in mutual conflict as they 'compete' in drawing strength from the same source.

In particular, as shown in Figure 2, we add a meta-argument on every support relationship in the original PRAF, and all the meta-arguments supported by the same page attack each other. While the 'regular' arguments still represent the pages, these new meta-arguments correspond to the links between them. This increases the expressivity of the representation, allowing in particular attacks between the meta-arguments corresponding to links from the same page in order to describe the fact that they 'compete' for conveying strength, as mentioned above, and therefore the more links originating from the same page, the lower the strength transferred through each of them.

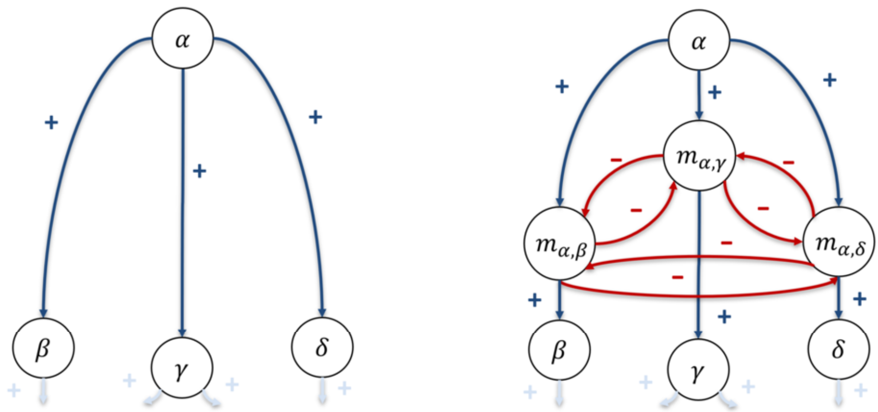

(a) PRAF (b) MPRAF

Figure 2. Example of a transformation from a PRAF to an MPRAF.

Definition 5. *Given a PRAF $PR = \langle \mathcal{X}, \emptyset, \mathcal{R}^+, \tau \rangle$, the* PageRank Meta-Argumentation Framework *(MPRAF) derived from PR is a QBAF $\langle \mathcal{X} \cup \mathcal{M}, \widehat{\mathcal{R}}^-, \widehat{\mathcal{R}}^+, \widehat{\tau} \rangle$, where:*

- $\mathcal{M} = \{m_{\alpha,\beta} : (\alpha, \beta) \in \mathcal{R}^+\}$ *is the set of meta-arguments*
- $\widehat{\mathcal{R}}^+ = \{(\alpha, m_{\alpha,\beta}), (m_{\alpha,\beta}, \beta) : \alpha, \beta \in \mathcal{X}, m_{\alpha,\beta} \in \mathcal{M}\}$ *is the set of supports*
- $\widehat{\mathcal{R}}^- = \{(m_{\alpha,\beta}, m_{\alpha,\gamma}) \in \mathcal{M} \times \mathcal{M} : (\alpha, \beta), (\alpha, \gamma) \in \mathcal{R}^+\}$ *is the set of attackers*
- $\widehat{\tau} : \mathcal{X} \cup \mathcal{M} \mapsto \widehat{\mathbb{I}} = [0, 1[$ *is the base score defined as the function:*

$$\widehat{\tau}(\alpha) = \begin{cases} 0 & \text{if } \alpha \in \mathcal{M} \\ \frac{1-d}{|\mathcal{X}|} & \text{if } \alpha \in \mathcal{X} \end{cases}$$

Figure 2 illustrates the transformation of a PRAF into an MPRAF: the supports go from a 'regular' argument to another through an intermediate meta-argument. The following remarks illustrate some of the properties of MPRAFs $\langle \mathcal{X} \cup \mathcal{M}, \widehat{\mathcal{R}}^-, \widehat{\mathcal{R}}^+, \widehat{\tau} \rangle$.

Remark 2. *For any $\alpha \in \mathcal{X}$, $\widehat{\mathcal{R}}^-(\alpha) = \emptyset$.*

Remark 3. *For any $m_{\alpha,\beta} \in \mathcal{M}$, $\exists! \alpha \in \widehat{\mathcal{R}}^+(m_{\alpha,\beta})$, $\exists! \beta \in \widehat{\mathfrak{R}}^+(m_{\alpha,\beta})$, $\alpha \in \mathcal{X}$ and $\beta \in \mathcal{X}$.*

Remark 4. *For any $m_{\alpha,\beta} \in \mathcal{M}$, $|\widehat{\mathcal{R}}^-(m_{\alpha,\beta})| + 1 = |\mathfrak{R}^+(\alpha)| = |\widehat{\mathfrak{R}}^+(\alpha)|$.*

Remark 5. *For any $\alpha \in \mathcal{X}$ where $\exists!\ m_{\alpha,\beta} : (\alpha, m_{\alpha,\beta}) \in \widehat{\mathcal{R}}^+$, $\widehat{\mathcal{R}}^-(m_{\alpha,\beta}) = \emptyset$.*

With reference to MPRAFs, we now define a gradual semantics $\widehat{\sigma}$, whose outcomes on 'regular' arguments coincide with the score produced by PR, as proved in Thm. 1.

Definition 6. *The* Meta-PageRank semantics *(M-PR) is a gradual semantics $\widehat{\sigma} : \mathcal{X} \cup \mathcal{M} \mapsto \widehat{\mathbb{I}}$ such that:*

$$\widehat{\sigma}(\alpha) = \widehat{\tau}(\alpha) + \sqrt{d} \cdot \frac{\sum_{\beta \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\alpha)| + 1} \quad \forall \alpha \in \mathcal{X} \cup \mathcal{M}$$

We now prove that, given a PRAF and corresponding MPRAF, for any $\alpha \in \mathcal{X}$, the strength $\widehat{\sigma}(\alpha)$ according to Def. 6 is the same as the strength $\sigma(\alpha)$ according to Def. 4, i.e. to the PR score.

Theorem 1. *Given a PRAF $\langle \mathcal{X}, \emptyset, \mathcal{R}^+, \tau \rangle$, denoted as PR, and the corresponding MPRAF $\langle \mathcal{X} \cup \mathcal{M}, \widehat{\mathcal{R}}^-, \widehat{\mathcal{R}}^+, \widehat{\tau} \rangle$, denoted as $\widehat{PR}$, with the semantics σ for PR and $\widehat{\sigma}$ for $\widehat{PR}$, for any argument $\alpha \in \mathcal{X}$ it holds that $\sigma(\alpha) = \widehat{\sigma}(\alpha)$.*

Proof. By Def. 6, $\widehat{\sigma}(\alpha) = \frac{1-d}{|\mathcal{X}|} + \sqrt{d} \cdot \frac{\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma)}{|\widehat{\mathcal{R}}^-(\alpha)|+1}$. By hypothesis $\alpha \in \mathcal{X}$, thus if $\gamma \in \widehat{\mathcal{R}}^+(\alpha)$ then $\gamma \in \mathcal{M}$, so we can rewrite γ as $m_{\beta,\alpha}$ where $\beta \in \mathcal{R}^+(\alpha)$. By the same hypothesis, we can derive, by Rem. 2, that $|\widehat{\mathcal{R}}^-(\alpha)| = 0$. This means that $\widehat{\sigma}(\alpha)$ can be rewritten as $\frac{1-d}{|\mathcal{X}|} + \sqrt{d} \cdot \sum_{m_{\beta,\alpha} \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(m_{\beta,\alpha})$. Expliciting $\widehat{\sigma}(m_{\beta,\alpha})$ by Def. 6 and recalling that, by Def. 5, $\tau(m_{\beta,\alpha}) = 0$ because $m_{\beta,\alpha}$ is a meta-argument, $\widehat{\sigma}(\alpha) = \frac{1-d}{|\mathcal{X}|} + \sqrt{d} \cdot \sum_{m_{\beta,\alpha} \in \widehat{\mathcal{R}}^+(\alpha)} \left(\sqrt{d} \cdot \frac{\sum_{\beta \in \widehat{\mathcal{R}}^+(m_{\beta,\alpha})} \widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(m_{\beta,\alpha})|+1} \right)$. We recall that, by Rem. 3, $\exists! \beta : \beta \in \widehat{\mathcal{R}}^+(m_{\beta,\alpha})$ because $m_{\beta,\alpha} \in \mathcal{M}$. Furthermore, we know by Rem. 4 that $|\widehat{\mathcal{R}}^-(m_{\beta,\alpha})| + 1 = |\mathfrak{R}^+(\beta)|$. Thus, $\widehat{\sigma}(\alpha) = \frac{1-d}{|\mathcal{X}|} + d \cdot \sum_{m_{\beta,\alpha} \in \widehat{\mathcal{R}}^+(\alpha)} \frac{\widehat{\sigma}(\beta)}{|\mathfrak{R}^+(\beta)|}$. This is equivalent to $\widehat{\sigma}(\alpha) = \frac{1-d}{|\mathcal{X}|} + d \cdot \sum_{\beta \in \mathcal{R}^+(\alpha)} \frac{\widehat{\sigma}(\beta)}{|\mathfrak{R}^+(\beta)|} = \sigma(\alpha)$. □

Lemma 2 proves that the codomain of $\widehat{\sigma}$ is $\widehat{\mathbb{I}}$.

Lemma 2. *The codomain of $\widehat{\sigma}$ on an MPRAF $\langle \mathcal{X} \cup \mathcal{M}, \widehat{\mathcal{R}}^-, \widehat{\mathcal{R}}^+, \widehat{\tau} \rangle$ is $\widehat{\mathbb{I}} =]0, 1]$. Moreover, for any $\alpha \in \mathcal{X} \cup \mathcal{M}$, if $\alpha \in \mathcal{X}$ then $\widehat{\sigma}(\alpha) \geq \frac{1-d}{|\mathcal{X}|}$, otherwise $\widehat{\sigma}(\alpha) > 0$.*

Proof. By Def. 6, $\widehat{\sigma}(\alpha)$ is the sum of $\widehat{\tau}(\alpha)$ and positive values. Hence if $\alpha \in \mathcal{X}$ then $\widehat{\sigma}(\alpha) \geq \frac{1-d}{|\mathcal{X}|} > 0$. Otherwise, if $\alpha \in \mathcal{M}$ then, by Defs. 5 and 6, $\widehat{\sigma}(\alpha) = \sqrt{d} \cdot \frac{\sum_{\beta \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\alpha)|+1} \geq \sqrt{d} \cdot \sum_{\beta \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\beta)$, and since $\beta \in \mathcal{X}$ then $\widehat{\sigma}(\beta) > 0 \quad \forall \beta$, hence $\widehat{\sigma}(\alpha) > 0$. By Theorem 1 and by Lem. 1, we have that if $\alpha \in \mathcal{X}$ then $\widehat{\sigma}(\alpha) \leq 1$. Otherwise, if $\alpha \in \mathcal{M}$ then, by Rem. 3, $\widehat{\mathcal{R}}^+(\alpha) = \{\beta\}$ and $\beta \in \mathcal{X}$, hence by Def. 6, $\widehat{\sigma}(\alpha) = \sqrt{d} \cdot \frac{\widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\alpha)|+1} \leq 1$. □

The next proposition sheds light on the intuition behind our MPRAFs, in that the support from non-meta-arguments is partitioned among the meta-arguments. Meta-arguments supported by the same 'regular' argument all have the same strength since according to the random surfer model the probability of clicking on links is uniform.

Proposition 5. *In an MPRAF $\langle \mathcal{X} \cup \mathcal{M}, \widehat{\mathcal{R}}^-, \widehat{\mathcal{R}}^+, \widehat{\tau} \rangle$, if a meta-argument $\alpha \in \mathcal{M}$ has attackers then $\widehat{\sigma}(\alpha) = \widehat{\sigma}(\gamma), \ \forall \gamma \in \widehat{\mathcal{R}}^-(\alpha)$.*

Proof. By Def. 5, $\forall \gamma \in \widehat{\mathcal{R}}^-(\alpha) \quad \gamma \in \mathcal{M}$ and by Def. 5 and Rem. 3 $\forall \gamma \in \widehat{\mathcal{R}}^-(\alpha) \quad \widehat{\mathcal{R}}^+(\alpha) = \widehat{\mathcal{R}}^+(\gamma) = \{\beta\}$ where $\beta \in \mathcal{X}$ is the single supporter of α. By Def. 6, $\widehat{\sigma}(\alpha) = \widehat{\tau}(\alpha) + \sqrt{d} \cdot \frac{\sum_{\beta \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\alpha)|+1}$, and by Def. 5 and Rem. 3, $\widehat{\sigma}(\alpha) = \sqrt{d} \cdot \frac{\widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\alpha)|+1}$, and the same is true for any $\gamma \in \widehat{\mathcal{R}}^-(\alpha)$: $\widehat{\sigma}(\gamma) = \sqrt{d} \cdot \frac{\widehat{\sigma}(\beta)}{|\widehat{\mathcal{R}}^-(\gamma)|+1}$. By construction α and the elements of $\widehat{\mathcal{R}}^-(\alpha)$ all attack each other, thus $|\widehat{\mathcal{R}}^-(\alpha)| = |\widehat{\mathcal{R}}^-(\gamma)| \ \forall \gamma \in \widehat{\mathcal{R}}^-(\alpha)$, and the result follows. □

We now assess this framework and semantics with respect to the desirable properties.

Proposition 6. *$\widehat{\sigma}$ satisfies GP1, GP4, GP5, GP6, GP8, GP9 and GP11.*

Proof. GP1: by Def. 6, if $\widehat{\mathcal{R}}^+(\alpha) = \emptyset$ and $\widehat{\mathcal{R}}^-(\alpha) = \emptyset$ then the second term of the sum is always 0, therefore $\sigma(\alpha) = \tau(\alpha)$. GP4 holds because the GP's preconditions cannot be verified: by Lem. 2, $\forall \alpha \in \mathcal{X} \quad \widehat{\sigma}(\alpha) \geq \widehat{\tau}(\alpha)$. GP5: by Def. 6, $\widehat{\sigma}(\alpha) > \widehat{\tau}(\alpha)$ iff $\sum_{\beta \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\beta) > 0$. Thus, it must be the case that $\exists \beta \in \widehat{\mathcal{R}}^+(\alpha) : \widehat{\sigma}(\beta) > 0$, therefore $\widehat{\mathcal{R}}^+(\alpha) \neq \emptyset$ GP6: follows directly from Def. 6. GP8: if $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^-(\beta)$ then $|\widehat{\mathcal{R}^-_\sigma}(\alpha)| = |\widehat{\mathcal{R}^-_\sigma}(\beta)|$ and if $\widehat{\mathcal{R}^+_\sigma}(\alpha) \subsetneq \widehat{\mathcal{R}^+_\sigma}(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) < \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. The result follows from Def. 6. GP9: if $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^-(\beta)$ then $|\widehat{\mathcal{R}^-_\sigma}(\alpha)| = |\widehat{\mathcal{R}^-_\sigma}(\beta)|$ and if $\widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^+(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) = \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. The result follows from Def. 6. GP11: if $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^-(\beta)$ then $|\widehat{\mathcal{R}^-_\sigma}(\alpha)| = |\widehat{\mathcal{R}^-_\sigma}(\beta)|$ and if $\widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{>} \widehat{\mathcal{R}}^+(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) > \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. The result follows from Def. 6. □

Proposition 7. *$\langle \widehat{PR}, \widehat{\sigma} \rangle$ is (not strictly) balanced and thus satisfies GP1 to GP3.*

Proof. Point 1: (A) If $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^+(\alpha) = \emptyset$ then the result follows by Def. 6. (B) Otherwise, if $\widehat{\mathcal{R}}^-(\alpha) \neq \emptyset$ then $\alpha \in \mathcal{M}$ and thus it has a single supporter β. There are two possible scenarios (B.i) $\exists! \gamma \in \mathcal{M} : (\beta, \alpha), (\beta, \gamma) \in \widehat{\mathcal{R}}^+$, then $\{\beta\} = \widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{>} \widehat{\mathcal{R}}^-(\alpha) = \{\gamma\}$ (which contradicts the hypothesis) because by Def. 6 $\widehat{\sigma}(\alpha) = \widehat{\sigma}(\gamma) < \widehat{\sigma}(\beta)$ (B.ii) $\exists_{>1} \gamma_1, ..., \gamma_n \in \mathcal{M} : (\beta, \alpha), (\beta, \gamma_1), ..., (\beta, \gamma_n) \in \widehat{\mathcal{R}}^+$, hence $|\widehat{\mathcal{R}}^-(\alpha)| > 1$, therefore it cannot hold that $\{\gamma_1, ..., \gamma_n\} = \widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^+(\alpha) = \{\beta\}$ (which contradicts the hypothesis), and by Def. 6 it holds again $\widehat{\sigma}(\alpha) = \widehat{\sigma}(\gamma_1) = ... = \widehat{\sigma}(\gamma_n) < \widehat{\sigma}(\beta)$, hence it cannot exists any injective mapping $f : \widehat{\mathcal{R}}^-(\alpha) \to \widehat{\mathcal{R}}^+(\alpha) : \forall \alpha \in \widehat{\mathcal{R}}^-(\alpha), \sigma(f(\alpha)) \geq \sigma(\alpha)$, and thus there is no strength-equivalency relationship between $\widehat{\mathcal{R}}^-(\alpha)$ and $\widehat{\mathcal{R}}^+(\alpha)$. Point 2. For $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{>} \widehat{\mathcal{R}}^+(\alpha)$ to hold $\widehat{\mathcal{R}}^-(\alpha) \neq \emptyset$, thus $\alpha \in \mathcal{M}$. Hence, we are in the same situation of (B) in the proof of Point 1, and therefore the precondition cannot hold and the result follows. Point 3. By Lem. 2, $\widehat{\sigma}(\alpha) > 0$ and if $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{<} \widehat{\mathcal{R}}^+(\alpha)$ then $\widehat{\mathcal{R}}^+(\alpha) \neq \emptyset$. Hence by Def. 6, $\widehat{\sigma}(\alpha) > \widehat{\tau}(\alpha)$. Point 4 holds because $\nexists \alpha : \widehat{\sigma}(\alpha) < \widehat{\tau}(\alpha)$. Point 5 does not hold. For example, consider the framework in Figure 2.b and in particular $m_{\alpha,\gamma} \in \mathcal{M}$ that it is supported by $\alpha \in \mathcal{X}$ and attacked by $m_{\alpha,\beta}, m_{\alpha,\delta} \in \mathcal{M}$. By Def. 5 and Lem. 2, we have that $\widehat{\sigma}(m_{\alpha,\gamma}) \leq \widehat{\sigma}(\alpha)$ and $\widehat{\sigma}(m_{\alpha,\gamma}) = \widehat{\sigma}(m_{\alpha,\beta}) = \widehat{\sigma}(m_{\alpha,\delta}) > 0$. Hence, $\widehat{\sigma}(m_{\alpha,\gamma}) > \widehat{\tau}(m_{\alpha,\gamma})$, but $\widehat{\mathcal{R}}^+(m_{\alpha,\gamma}) \stackrel{\sigma}{\ngeq} \widehat{\mathcal{R}}^-(m_{\alpha,\gamma})$ because no injective mapping exists from $\widehat{\mathcal{R}}^-(m_{\alpha,\gamma})$ to $\widehat{\mathcal{R}}^+(m_{\alpha,\gamma})$. Thus $\widehat{\mathcal{R}}^+(m_{\alpha,\gamma}) \stackrel{\sigma}{\not>} \widehat{\mathcal{R}}^-(m_{\alpha,\gamma})$. □

Proposition 8. *$\langle \widehat{PR}, \widehat{\sigma} \rangle$ is strictly monotonic and thus satisfies GP6 to GP11.*

Proof. Point 1: if $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^-(\beta)$ then $|\widehat{\mathcal{R}}^-(\alpha)| = |\widehat{\mathcal{R}}^-(\beta)|$ and if $\widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^+(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) = \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. The result follows from Def. 6. Point 3: if $\alpha, \beta \in \mathcal{X}$ then $\widehat{\tau}(\alpha) = \widehat{\tau}(\beta)$ and $\widehat{\mathcal{R}}^-(\beta) \stackrel{\sigma}{=} \widehat{\mathcal{R}}^-(\alpha) = \emptyset$, hence $|\widehat{\mathcal{R}}^-(\alpha)| = |\widehat{\mathcal{R}}^-(\beta)|$. If $\widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{<} \widehat{\mathcal{R}}^+(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) < \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. Thus, by Def. 6, $\widehat{\sigma}(\alpha) < \widehat{\sigma}(\beta)$. If $\alpha \in \mathcal{M}$ and $\beta \in \mathcal{X}$ then $\widehat{\tau}(\alpha) < \widehat{\tau}(\beta)$ and $\widehat{\mathcal{R}}^-(\beta) = \emptyset$. If $\widehat{\mathcal{R}}^-(\alpha) \stackrel{\sigma}{\geq} \emptyset$ then $|\widehat{\mathcal{R}}^-(\alpha)| \geq |\widehat{\mathcal{R}}^-(\beta)| = 0$. If $\widehat{\mathcal{R}}^+(\alpha) \stackrel{\sigma}{\leq} \widehat{\mathcal{R}}^+(\beta)$ then

Table 1. GPs and principles (**B**alance, **S**trict **B**alance, **M**onotonicity, **S**trict **M**onotonicity) satisfied by σ, $\langle PR, \sigma \rangle$, $\widehat{\sigma}$ and $\langle \widehat{PR, \widehat{\sigma}} \rangle$, where ✓ and $\times$ denote property satisfied and not satisfied, resp.

	GP1	GP2	GP3	GP4	GP5	GP6	GP7	GP8	GP9	GP10	GP11	B	SB	M	SM
σ	✓	$\times$	✓	✓	✓	$\times$	$\times$	✓	✓	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$
$\langle PR, \sigma \rangle$	✓	✓	✓	✓	✓	$\times$	✓	✓	✓	✓	$\times$	✓	✓	$\times$	$\times$
$\widehat{\sigma}$	✓	$\times$	$\times$	✓	✓	✓	$\times$	✓	✓	$\times$	✓	$\times$	$\times$	$\times$	$\times$
$\langle \widehat{PR, \widehat{\sigma}} \rangle$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	$\times$	✓	✓

$\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) \leq \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. Thus, by Def. 6, $\widehat{\sigma}(\alpha) < \widehat{\sigma}(\beta)$. If $\alpha, \beta \in \mathcal{M}$ then $\widehat{\tau}(\alpha) = \widehat{\tau}(\beta)$. If $\widehat{\mathcal{R}}^-(\beta) \overset{\sigma}{\leq} \widehat{\mathcal{R}}^-(\alpha)$ then $|\widehat{\mathcal{R}}^-(\alpha)| \geq |\widehat{\mathcal{R}}^-(\beta)|$. If $\widehat{\mathcal{R}}^+(\alpha) \overset{\sigma}{\leq} \widehat{\mathcal{R}}^+(\beta)$ then $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) \leq \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$. Hence, by Def. 6, $\widehat{\sigma}(\alpha) \leq \widehat{\sigma}(\beta)$. For $\mathcal{ST}(\beta) \not\preceq \mathcal{ST}(\alpha)$ to hold, either:

- $\widehat{\mathcal{R}}^-(\beta) \overset{\sigma}{<} \widehat{\mathcal{R}}^-(\alpha)$ and $\widehat{\mathcal{R}}^+(\alpha) \overset{\sigma}{=} \widehat{\mathcal{R}}^+(\beta)$, or
- $\widehat{\mathcal{R}}^-(\beta) \overset{\sigma}{=} \widehat{\mathcal{R}}^-(\alpha)$ and $\widehat{\mathcal{R}}^+(\alpha) \overset{\sigma}{<} \widehat{\mathcal{R}}^+(\beta)$, or
- $\widehat{\mathcal{R}}^-(\beta) \overset{\sigma}{<} \widehat{\mathcal{R}}^-(\alpha)$ and $\widehat{\mathcal{R}}^+(\alpha) \overset{\sigma}{<} \widehat{\mathcal{R}}^+(\beta)$.

In the first case, by construction of the framework PR, $|\widehat{\mathcal{R}}^-(\alpha)| < |\widehat{\mathcal{R}}^-(\beta)|$, thus $\widehat{\sigma}(\alpha) < \widehat{\sigma}(\beta)$. In the second case, $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) < \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$, thus $\widehat{\sigma}(\alpha) < \widehat{\sigma}(\beta)$. In the third case, $\sum_{\gamma \in \widehat{\mathcal{R}}^+(\alpha)} \widehat{\sigma}(\gamma) < \sum_{\gamma \in \widehat{\mathcal{R}}^+(\beta)} \widehat{\sigma}(\gamma)$ and $|\widehat{\mathcal{R}}^-(\alpha)| \leq |\widehat{\mathcal{R}}^-(\beta)|$, thus $\widehat{\sigma}(\alpha) < \widehat{\sigma}(\beta)$. Point 3 implies Point 2, thus the result follows. $\square$

We have thus proven that, through MPRAF, in exchange for a little structural addition, it is possible to ensure equivalence with PR while at the same time satisfying more desirable properties from an argumentation semantics perspective. The value of the proposed approach is not purely theoretical however, as we discuss in next section.

5. Towards better PR explanations

As shown in Table 1, the M-PR semantics applied on an MPRAF satisfies almost all the desirable properties outlined in Section 3, including in particular *monotonicity*. This means that, from a dialectical viewpoint, the strength of an argument depends exclusively on its intrinsic strength, the reasons supporting it and the reasons against it, and any strengthening/weakening of these will affect the argument's strength intuitively. The satisfaction of monotonicity is achieved through the role ascribed to meta-arguments and is a key factor for exploiting MPRAFs for practical application, such as the generation of explanations of the PR score of a page. In this scenario, *monotonicity* is clearly a crucial factor because it allows a user to identify direct dependencies between the variations of the strength of arguments according to the attacks and supports linking them in the graph structure of the MPRAF. For this reason, MPRAFs are able to provide the end user a better understanding of the factors determining the PR score of a page, i.e. they support answering questions like "Which incoming links (and thus pages) contribute the most to the score of this page?".

To provide some preliminary empirical support to this claim, we ran some experiments on the Wikipedia dataset from Wikipedia Dumps consisting of 965,748 pages and 7,388,700 links, with an average link density of 7.65 links per page. While discussing more extensively our experiments is beyond the scope of this paper, we provide a concrete example of the explanatory advantages achievable with MPRAFs.

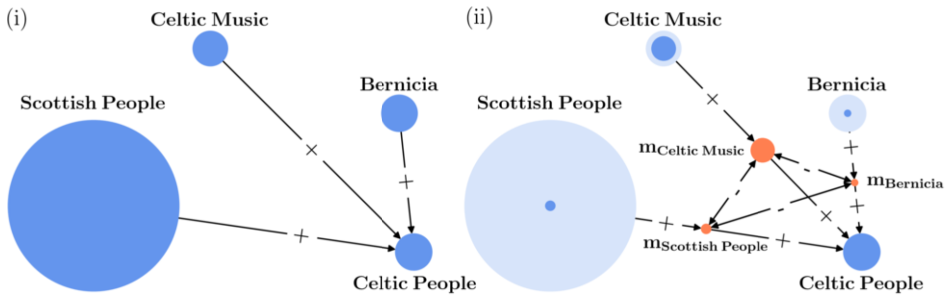

Figure 3. Excerpt of the PRAF (i) and MPRAF (ii) for the article *Celtic people* in the *simple* version of *Wikipedia* including the article and its direct supporters. Each bubble represents an argument and its size is proportional the the strength of the argument. In (ii) the opaque bubbles highlight the actual contribution of an argument to the *Celtic People* page, derived from the strengths of the corresponding meta-arguments.

Consider first Fig. 3.i, showing a magnification of the weighted view of the pages contributing to the score of the article *Celtic People*, with each page score represented by the size of the relevant bubble. Looking at this figure a user might (erroneously) deduce that the score of *Celtic People* is mostly determined by *Scottish People*, which is actually not the case (due to the high number of outgoing links from *Scottish People*). To realize this a user should both have a deeper understanding of PR's functioning and be shown a larger part of the graph, including all the pages linked by *Celtic People*'s supporters.

This undesirable overload is avoided by the MPRAF-based representation in Figure 3.ii. Here the meta-arguments show directly the actual support flowing from the supporters, and the user can appreciate that *Celtic Music* is the article providing most support to *Celtic People*. Besides better supporting direct explanations, the MPRAF-based representation appears to enable answering other kinds of user queries, like counterfactual questions of the kind: 'What would happen if a given link is suppressed?' A wider investigation on MPRAF-based explanations for PR outcomes is planned for future work.

6. Conclusions

Towards the more general goal of investigating connections between PageRank and argument evaluation, we have introduced a novel approach capable of reconstructing PR as a gradual argumentation semantics of a suitably defined bipolar argumentation framework, while ensuring the satisfaction of a set of generally desirable properties. We have then given an example of the practical yields of this theoretical achievement, concerning the generation of better explanations of PR scores to end users.

To the best of our knowledge, the investigation of the relationships between PR and argumentation semantics has not been previously considered in the literature. The work in [15] explores the application of PR to rank the relevance of arguments available on the web to support or attack a given stance. This is an interesting but different goal: in [15] PR is not related to any semantics notion and the links have a different meaning, relating the conclusion of an argument with the premises of another one. On a different but related line, some works, e.g. [16], have explored connections between argumentation semantics and matrix representations from network theory, whose relationships with our approach are worth future investigation.

Our proposal can be extended in several directions. On one hand, the investigation of PR-inspired gradual semantics for various kinds of argumentation frameworks could be pursued. In this respect it would be interesting to consider *weighted* versions of PR where a node strength can be distributed unevenly to its children and more generally the variants of PR considered in various domains [6]. On the other hand, one can notice that PR is essentially a mechanism to produce a score based on a relation of support, but it could be considered that in several domains where PR is applied, also other relations, in particular attack could be relevant for a proper scoring. Also, in the web domain, one could argue that the absence of a link from one page to another (where this link could instead be expected according to some criterion) could be interpreted as an attack diminishing the relevance of the non-linked page. Given the strong tradition on attack-based and bipolar evaluations in argumentation semantics, this suggests that the study of argumentation-inspired variants of PR may also represent a fruitful research direction.

References

[1] Page L, Brin S, Motwani R, Winograd T. The PageRank Citation Ranking: Bringing Order to the Web. World Wide Web Internet And Web Information Systems. 1998;54(1999-66):1–17.

[2] Ma N, Guan J, Zhao Y. Bringing PageRank to the citation analysis. Information Processing and Management. 2008 3;44(2):800–810.

[3] Gori M, Pucci A. ItemRank: A Random-Walk Based Scoring Algorithm for Recommender Engines. In: Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI); 2007. p. 2766–2771.

[4] Hudelson M, Mooney BL, Clark AE. Determining polyhedral arrangements of atoms using PageRank. Journal of Mathematical Chemistry. 2012 9;50(9):2342–2350.

[5] Morrison JL, Breitling R, Higham DJ, Gilbert DR. GeneRank: Using search engine technology for the analysis of microarray experiments. BMC Bioinformatics. 2005 sep;6(1):233.

[6] Gleich DF. PageRank beyond the web. SIAM Review. 2015;57(3):321–363.

[7] Altman A, Tennenholtz M. Ranking systems: the PageRank axioms. In: Proc. 6th ACM Conf. on Electronic Commerce (EC); 2005. p. 1–8.

[8] Dung PM. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. Artificial Intelligence. 1995;77(2):321–358.

[9] Cayrol C, Lagasquie-Schiex MC. On the Acceptability of Arguments in Bipolar Argumentation Frameworks. In: Proc. of the 8th European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU); 2005. p. 378–389.

[10] Baroni P, Rago A, Toni F. From fine-grained properties to broad principles for gradual argumentation: A principled spectrum. Int Journal Approximate Reasoning. 2019;105:252–286.

[11] Baroni P, Caminada M, Giacomin M. An introduction to argumentation semantics. Knowledge Engineering Review. 2011;26(4):365–410.

[12] Cayrol C, Lagasquie-Schiex MC. Graduality in Argumentation. Journal of Artificial Intelligence Research. 2005;23:245–297.

[13] Amgoud L, Ben-Naim J. Evaluation of Arguments from Support Relations: Axioms and Semantics. In: Proc. of the 25th Int. Joint Conf. on Artificial Intelligence (IJCAI); 2016. p. 900–906.

[14] Langville AN, Meyer CD. Deeper inside PageRank. Internet Mathematics. Internet Mathematics. 2004;1(3):335–380.

[15] Wachsmuth H, Stein B, Ajjour Y. "PageRank" for Argument Relevance. In: Proc. of the 15th Conf. of the European Chpt. of the Association for Computational Linguistics (EACL); 2017. p. 1117–1127.

[16] Corea C, Thimm M. Using Matrix Exponentials for Abstract Argumentation. In: Proc. of the 1st Int. Workshop on Systems and Algorithms for Formal Argumentation (SAFA); 2016. p. 10–21.