ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230515

Graduality in Probabilistic Argumentation Frameworks

Jeroen Spaans[a] and Dragan Doder[a]

[a]Utrecht University
ORCiD ID: Jeroen Spaans https://orcid.org/0000-0001-7027-8102,
Dragan Doder https://orcid.org/0000-0003-0067-3654

Abstract. Gradual semantics are methods that evaluate overall strengths of individual arguments in graphs. In this paper, we investigate gradual semantics for extended frameworks in which probabilities are used to quantify the uncertainty about arguments and attacks belonging to the graph. We define the likelihoods of an argument's possible strengths when facing uncertainty about the topology of the argumentation framework. We also define an approach to compare the strengths of arguments in this probabilistic setting. Finally, we propose a method to calculate the overall strength of each argument in the framework, and we evaluate this method against a set of principles.

1 Introduction

Within the last couple of decades, argumentation has emerged as a popular field in Artificial Intelligence [9, 12]. It has been shown to be useful in several domains, such as decision making [40], reasoning under inconsistency [13] and non-monotonic reasoning [29] and is applicable in the domains of law and medicine [9]. The underlying structure of formal models of abstract argumentation takes the form of directed graphs, whose nodes represent arguments and whose directed edges indicate attacks between attacks.

Two main classes of semantics were proposed to reason about such structures and to evaluate arguments in the graphs. *Extension-based* semantics are proposed with the goal of identifying jointly acceptable sets of arguments (extensions) based on specific properties within the graph [18]. The acceptability status of an argument is then derived from these extensions. The argument is *sceptically accepted* if it belongs to all extensions, *credulously accepted* if it belongs to some of the extensions, and *rejected* otherwise. On the other hand, *gradual semantics* [17] focus on individual arguments and quantify their strengths in graphs, using a richer scale (usually the unit interval of reals $[0, 1]$). They typically define strength of an argument to depend on strengths of its direct attackers. One intuitive difference between gradual and extension-based semantics is that, in the latter approach, the attack relation is used to destroy its target (two conflicting arguments cannot be in the same extension), while in ranking based semantics it is often used to only weaken its target. Examples of gradual semantics are h-Categoriser [13], Simple product semantics [27], Trust-based semantics[32], Iterative Schema [20], Max-based and Cardinality-based semantics [8]. Some of the approaches were adapted to frameworks where arguments and/or attacks have a base weight [5, 6, 7, 8], bipolar frameworks in which both support and attack relations are present in a graph [4, 34, 35], and weighted bipolar SETAFs where a set of arguments can attack or support a target together [39]. Gradual semantics are similar in spirit to *ranking semantics* [2, 16], which focus on the strength of arguments relative to that of other arguments, and return a preorder on arguments, thus ranking them from the strongest to the weakest ones. Obviously, each gradual semantics can be used to generate a ranking semantics (but not vice versa).

For many applications in which there is uncertainty about topology of the argumentation graph, simple attack frameworks appear too simple for convenient modelling of those aspects of an argumentation problem. There are different scenarios in which uncertainty about presence or arguments and attacks in a graph arises naturally: at times, ambiguities in the language used for presenting an argument, or the presentation of arguments with incomplete premises or claims may lead to uncertainty about the correct interpretation of attacks between arguments [23]; other times, arguments are presented with explicit uncertainty in their claims [23]; an audience to some argumentation is often unsure of the exact set of arguments being put forward, and a participant in in some argumentation may be unsure which arguments the audience has in mind [25]. In order to handle these uncertainties, Li, Oren and Norman [28] proposed Probabilistic Argumentation Frameworks which augment argumentation graphs with probabilities. In this approach, named the *constellations approach* by Hunter[1] [22], probability values are added to the arguments and attacks, allowing for the modelling of uncertainty in which elements should be present in the argumentation framework. In the constellations approach, a considered extension semantics is used to determine the probability of an argument being (credulously/sceptically) accepted. This approach to probabilistic argumentation has been extensively investigated from an extension-based semantic point of view [21, 22, 23, 25, 33, 14, 30, 24, 31, 15, 19], but never from the perspective of gradual semantics. Therefore, the current results on the constellations approach are well suited to answering the question of probability of (joint) acceptance of (sets of) arguments, but not the questions about probability of strength of an individual argument, or the question which argument is stronger in a probabilistic setting. Such questions are instead a particularly good match for applications of gradual semantics to graphs augmented with probabilities.

In this paper we study gradual semantics in probabilistic argumentation frameworks, following the constellations approach. This allows us to define semantics that return probabilities of an argument's acceptance with respect to any strength threshold, providing a richer scale of acceptability statuses than would be possible using Dung semantics (Section 3). In addition, we extend the approach to calculate

[1] In the same paper, the author introduces the notion of the *epistemic approach* to probabilistic argumentation, which can be used to represent the degree to which an argument is believed [22, 26, 37].

the probability of an argument being stronger than another argument. Moreover, we extend that notion of probabilistic ranking by defining the probabilities of certain *ranking queries* of interest. For example, we can calculate probability that the argument a is stronger then either b or c (Section 4). We investigate formal properties of both approaches and connections between them.

We also investigate the challenging problem of defining semantics that assign unique overall strength to each argument in a probabilistic argumentation framework (Section 5). Following the constellations approach, we propose desirable principles for such semantics[2], inspired by existing principles for gradual semantics from the literature [5]. We investigate properties of those principles and show their compatibility. Finally, we propose the first family of such semantics, by providing a method for generalising gradual semantics for argumentation graphs to our probabilistic framework. We show that if the considered underlying gradual semantics satisfies existing principles [5], then its generalisation satisfies our novel principles.

2 Background

2.1 Gradual semantics for argumentation frameworks

Arguments can be defeasible and may conflict with one another. Argumentation graphs, introduced by Dung under the name argumentation frameworks [18], are directed graphs that model such conflict by representing arguments with nodes and attacks between arguments with edges. This formalism considers arguments and attacks as purely abstract entities, doing away entirely with features such as the structure and origin of arguments and the nature of attacks.

Definition 2.1 (Argumentation Graph) *An argumentation graph (AG) is an ordered pair $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, where $\mathcal{A}$ is a non-empty finite set of arguments and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ is an attack relation between arguments. Let $\mathtt{AG}$ denote the set of all argumentation graphs.*

As we explained in the introduction, gradual semantics (unlike extension-based ones) assign to each argument a unique overall strength value or acceptability degree considering the strengths of their attackers. This richer evaluative scale, where arguments with a higher acceptability degree are considered more acceptable, enables us to make comparisons between alternative arguments, even when both would be assigned the same acceptability status by a Dung semantics.

Following an approach already accepted by some authors from the field [8], for simplicity we use the unit interval of reals as the evaluative scale.

Definition 2.2 (Weighting) *[8] A weighting on a set $\mathcal{X}$ is a function from $\mathcal{X}$ to the interval* $[0, 1]$.

Now we can define gradual semantics in a formal way.

Definition 2.3 (Gradual Semantics) *[8] A semantics is a function $\mathbf{S}$ transforming any argumentation graph $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle \in \mathtt{AG}$ into a weighting $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}$ on $\mathcal{A}$ (i.e., $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}} : \mathcal{A} \to [0, 1]$). For any $a \in \mathcal{A}$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a)$ represents the strength of a.*

A well-studied gradual semantics that will be used as an example in this work is the h-categoriser, proposed by Besnard and Hunter [13]:

[2] Interestingly, there is another recent principle-based study of argument strength in probabilistic argumentation [36]. However, that paper considers probabilities that are assigned to arguments' structure in logic-based argumentation, and therefore it does not belong to the constellations approach.

Definition 2.4 *The h-categoriser is a gradual semantics* $\mathtt{Hbs}$ *s.t.* $\forall \mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle \in \mathtt{AG}$, $\forall a \in \mathcal{A}$,

$$\mathtt{Deg}^{\mathtt{Hbs}}_{\mathbf{G}}(a) = \frac{1}{1 + \sum_{b_i \ s.t.\ (b_i, a) \in \mathcal{R}} \mathtt{Deg}^{\mathtt{Hbs}}_{\mathbf{G}}(b_i)}.$$

In the literature on gradual semantics, several works have been devoted to development of *principles* that represent desirable formal properties of semantics, with the purpose to serve as a tool for analysis and comparison of gradual semantics [3, 5, 8, 10, 11].

Here we informally introduce some of them (the corresponding formal definitions are listed in Section 2.2:

Maximality states that a non-attacked argument will have maximal strength (i.e., the strength 1).

Resilience states that every argument will have strictly positive strength in any graph.

Weakening states that if an argument has an attacker with positive strength, then its own strength cannot be maximal.

2.2 Principles for gradual semantics

In what follows, we recall some principles from [5], adjusted to non-weighted graphs. Let $\mathbf{S}$ be a gradual semantics.

Anonymity: $\mathbf{S}$ satisfies *anonymity* iff for any AGs $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle$, for any isomorphism f from $\mathbf{G}$ to $\mathbf{G}'$, the following holds: $\forall\, a \in \mathcal{A}, \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(f(a))$.

Independence: $\mathbf{S}$ satisfies *independence* iff for any AGs $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ s.t $\mathcal{A} \cap \mathcal{A}' = \emptyset$, the following holds: $\forall\, a \in \mathcal{A},\ \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G} \oplus \mathbf{G}'}(a)$ where $\mathbf{G} \oplus \mathbf{G}' = \langle \mathcal{A} \cup \mathcal{A}', \mathcal{R} \cup \mathcal{R}' \rangle$.

Directionality: $\mathbf{S}$ satisfies *directionality* iff for any AGs $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{G}' = \langle \mathcal{A}, \mathcal{R}' \rangle$ s.t. $\mathcal{R}' = \mathcal{R} \cup \{(a, b)\}$ and $\forall r \in \mathcal{R}, P'_{\mathcal{R}}(r) = P_{\mathcal{R}}(r)$, it holds that: $\forall x \in \mathcal{A}$, if there is no path from b to x, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(f(a))$.

Equivalence: $\mathbf{S}$ satisfies equivalence iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if there exists a bijective function f from $\mathtt{Att}_{\mathbf{G}}(a)$ to $\mathtt{Att}_{\mathbf{G}}(b)$ s.t. $\forall x \in \mathtt{Att}_{\mathbf{G}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(x) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(f(x))$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b)$.

Maximality: $\mathbf{S}$ satisfies maximality iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a \in \mathcal{A}$, it holds that: if $\mathtt{Att}_{\mathbf{G}}(a) = \emptyset$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = 1$.

Neutrality: $\mathbf{S}$ satisfies neutrality iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, it holds that: if $\mathtt{Att}_{\mathbf{G}}(b) = \mathtt{Att}_{\mathbf{G}}(a) \cup \{x\}$ s.t. $x \in \mathcal{A} \setminus \mathtt{Att}_{\mathbf{G}}(a)$ and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(x) = 0$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b)$.

Weakening: $\mathbf{S}$ satisfies weakening iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a \in \mathcal{A}$, it holds that if $\mathtt{Att}_{\mathbf{G}}(a) \neq \emptyset$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) < 1$.

Proportionality: $\mathbf{S}$ satisfies proportionality iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if $\mathtt{Att}_{\mathbf{G}}(a) = \mathtt{Att}_{\mathbf{G}}(b)$, $w(a) > w(b)$, and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > 0$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b)$.

Resilience: $\mathbf{S}$ satisfies resilience iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a \in \mathcal{A}$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > 0$.

Reinforcement: $\mathbf{S}$ satisfies reinforcement iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > 0$, $\mathtt{Att}_{\mathbf{G}}(a) \setminus \mathtt{Att}_{\mathbf{G}}(b) = \{x\}$, $\mathtt{Att}_{\mathbf{G}}(b) \setminus \mathtt{Att}_{\mathbf{G}}(a) = \{y\}$, and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(y) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(x)$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b)$.

Counting: $\mathbf{S}$ satisfies counting iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$, it holds that: if $\mathtt{Att}_{\mathbf{G}}(b) = \mathtt{Att}_{\mathbf{G}}(a) \cup \{x\}$ with $x \notin \mathtt{Att}_{\mathbf{G}}(a)$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(x) > 0$, and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > 0$, then $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) > \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b)$.

Weakening Soundness: $\mathbf{S}$ satisfies weakening soundness iff for any AG $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a \in \mathcal{A}$, it holds that: if $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) < 1$, then $\exists b \in \mathtt{Att}_{\mathbf{G}}(a)$ s.t. $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b) > 0$.

2.3 Probabilistic argumentation frameworks

When using argumentation graphs and semantics to evaluate arguments, we assume that all relevant arguments and attacks are considered and no arguments or attacks in the graph are irrelevant. In reality, however, it is rarely clear which arguments and attacks apply and should thus be placed in the graph. We might, for example, run into natural language being imprecise such that it is unclear whether claims are contradictory or there might be multiple ways to model the premises and claims often left implicit in conversation. We might also be in a situation where we are unsure what arguments a dialogue partner or an audience has in their mind. At other times we may be dealing with explicit uncertainty, such as with utterances like 'I am 99% sure' or 'I may have seen Robin yesterday'.

Such uncertainty is captured by probabilistic argumentation frameworks: an extension of argumentation frameworks, originally proposed by Li, Oren, and Norman [28], where each argument and attack is given a certain likelihood of appearing in a graph.

Definition 2.5 (PrAF) *A probabilistic argumentation framework (PrAF) is a quadruple* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$*, where* $\langle \mathcal{A}, \mathcal{R} \rangle$ *is an argumentation graph and* $P_{\mathcal{A}} : \mathcal{A} \rightarrow (0, 1]$ *and* $P_{\mathcal{R}} : \mathcal{R} \rightarrow (0, 1]$ *associate likelihood values with arguments and attacks respectively.* $\mathtt{PrAF}$ *denotes the set of all probabilistic argumentation frameworks.*

In a PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, the argumentation graph $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ represents the set of all arguments and attacks that may potentially appear. An instantiated graph that may arise under the uncertainty we face thus contains a subset of the arguments and attacks in $\mathbf{G}$. We call the process of deriving such a graph from a PrAF induction, and name the graphs that result induced graphs (of the PrAF). The functions $P_{\mathcal{A}}$ and $P_{\mathcal{R}}$ represent the uncertainty in the arguments and attacks in $\mathbf{G}$. $P_{\mathcal{A}}$ gives the probability that an argument appears in a graph induced from $\mathbf{G}$ and $P_{\mathcal{R}}$ the conditional probability that an attack appears in an induced graph given that both arguments it relates appear in the graph. The ranges of both functions deliberately exclude 0, as any argument or attack with a zero probability is known never to appear and is thus redundant. The maximum value 1 either function may assign represents certainty that an argument appears in an induced graph or that an attack appears given that its origin and target do as well.

We say that an argument is *perfect* in a PrAF, if it is both non-attacked and its probability is 1.

Definition 2.6 (Induced Graph) *An argumentation graph* $\mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ *is induced from a probabilistic argumentation framework* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ *iff all of the following hold:*

- $\mathcal{A}' \subseteq \mathcal{A}$
- $\mathcal{R}' \subseteq \mathcal{R} \cap (\mathcal{A}' \times \mathcal{A}')$
- $\forall a \in \mathcal{A}$ *such that* $P_{\mathcal{A}}(a) = 1$, $a \in \mathcal{A}'$
- $\forall (a, b) \in \mathcal{R}$ *such that* $P_{\mathcal{R}}((a, b)) = 1$ *and* $a, b \in \mathcal{A}'$, $(a, b) \in \mathcal{R}'$

$\mathtt{I}(\mathbf{F})$ *denotes the set of all argumentation graphs that may be induced from a probabilistic argumentation framework* $\mathbf{F}$.

This definition differs from the original, proposed in [28], in the fourth bullet; in addition to the attack (a, b) being certain, the original definition requires that arguments a and b both have probability 1 instead of only requiring them to be present in the graph for (a, b) to be certainly present in the graph. This change eliminates any induced graphs that will receive probability 0 under the following Definition 2.7, where an attack with probability 1 is not present even though the arguments it connects are.

Since independence between arguments and attacks is assumed, the probability of some induced graph $\mathbf{G}$ being induced from a PrAF $\mathbf{F}$ can be computed using the joint probabilities of independent variables:

Definition 2.7 (Probability of an Induced Graph) *Given a probabilistic argumentation framework* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$*, the probability of some graph* $\mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ *being induced from* $\mathbf{F}$ *is:*

$$P^I_{\mathbf{F}}(\mathbf{G}') = \prod_{a \in \mathcal{A}'} P_{\mathcal{A}}(a) \prod_{a \in \mathcal{A} \setminus \mathcal{A}'} (1 - P_{\mathcal{A}}(a)) \prod_{r \in \mathcal{R}'} P_{\mathcal{R}}(r) \prod_{r \in \mathcal{R}\downarrow_{\mathcal{A}'} \setminus \mathcal{R}'} (1 - P_{\mathcal{R}}(r))$$

where $\mathcal{R} \downarrow_{\mathcal{A}'} = \{(a, b) | a, b \in \mathcal{A}'$ *and* $(a, b) \in \mathcal{R}\}$.

Through definitions 2.5, 2.6 and 2.7 it is easily shown that

$$\forall \mathbf{F} \in \mathtt{PrAF},\ \forall \mathbf{G} \in \mathtt{I}(\mathbf{F}),\ P^I_{\mathbf{F}}(\mathbf{G}) > 0 \tag{1}$$

The following proposition claims that $P^I_{\mathbf{F}}$ is a probability distribution over the induced graphs of a PrAF.

Proposition 2.8 *The sum of probabilities of all argumentation graphs that may be induced from an arbitrary PrAF* $\mathbf{F}$ *is* 1*. i.e.* $\sum_{\mathbf{G} \in \mathtt{I}(\mathbf{F})} P^I_{\mathbf{F}}(\mathbf{G}) = 1$.

This distribution is used in [28] to define the probability of a set of arguments $\mathcal{X}$ being (sceptically/credulously) accepted in a PrAF as the sum of probabilities of all induced graphs that contain $\mathcal{X}$ in (some/all of) their extensions.

What follows is a list of special notations used in the paper.

Notation 1 *Let* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ *be a PrAF,* $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ *be an AG, and* $a \in \mathcal{A}$*. We write* $a\mathcal{R}b$ *iff* $(a, b) \in \mathcal{R}$*.* $\mathtt{Att}_{\mathbf{F}}(a)$ *and* $\mathtt{Att}_{\mathbf{G}}(a)$ *denote the set of all attackers of* a *in* $\mathbf{F}$ *and* $\mathbf{G}$ *respectively (i.e.,* $\mathtt{Att}_{\mathbf{F}}(a) = \mathtt{Att}_{\mathbf{G}}(a) = \{b \in \mathcal{A} | b\mathcal{R}a\}$*). For* $a, b \in \mathcal{A}$*, we say there is a path from* b *to* a *if there exists a finite non-empty sequence* $\langle x_1, \ldots, x_n \rangle$ *of arguments* $x_i \in \mathcal{A}$ *s.t.* $x_1 = b, x_n = a$ *and* $\forall i < n, x_i \mathcal{R} x_{i+1}$*. For any* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle, \mathbf{F}' = \langle \mathcal{A}', P'_{\mathcal{A}}, \mathcal{R}', P'_{\mathcal{R}} \rangle \in \mathtt{PrAF}$ *s.t.* $\mathcal{A} \cap \mathcal{A}' = \emptyset$*,* $\mathbf{F} \oplus \mathbf{F}'$ *is the PrAF* $\langle \mathcal{A} \cup \mathcal{A}', P''_{\mathcal{A}}, \mathcal{R} \cup \mathcal{R}', P''_{\mathcal{R}} \rangle$ *where for any* $a \in \mathcal{A}$ *(respectively* $a \in \mathcal{A}'$*)* $P''_{\mathcal{A}}(a) = P_{\mathcal{A}}(a)$ *(respectively* $P''_{\mathcal{A}}(a) = P'_{\mathcal{A}}(a)$*) and for any* $r \in \mathcal{R}$ *(respectively* $r \in \mathcal{R}'$*)* $P''_{\mathcal{R}}(r) = P_{\mathcal{R}}(r)$ *(respectively* $P''_{\mathcal{R}}(r) = P'_{\mathcal{R}}(r)$*). For any* $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle, \mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle \in \mathtt{AG}$ *s.t.* $\mathcal{A} \cap \mathcal{A}' = \emptyset$*,* $\mathbf{G} \oplus \mathbf{G}'$ *is the AG* $\langle \mathcal{A} \cup \mathcal{A}', \mathcal{R} \cup \mathcal{R}' \rangle$*. For any* $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ *and* $S \subseteq \mathcal{A}$*,* $\mathbf{F}_{|S} = \langle S, (P_{\mathcal{A}})_{|S}, \mathcal{R}_{|S \times S}, (P_{\mathcal{R}})_{|(\mathcal{R}_{|S \times S})} \rangle$.

3 Acceptability of arguments

In the work of Li, Oren, and Norman [28], Definition 2.7 is used to determine the probability that some set of arguments is acceptable under a given Dung semantics by adding together the probabilities of those induced graphs where the set is acceptable. This section explores the possibility of determining that probability under a gradual semantics instead.

Where Dung semantics determine the acceptability of an argument directly, without explicitly considering the overall strength of the argument, gradual semantics may be used to determine acceptability indirectly through the strengths they assign arguments. An approach to deriving argument acceptability from the strengths assigned to arguments by a gradual semantics is proposed in [1]; we may simply accept those arguments whose strength meets or exceeds a threshold we choose.

Definition 3.1 (Acceptability Under Gradual Semantics) *[1]* *Given an argumentation graph $\mathbf{G} = \langle \mathcal{A}, \mathcal{R} \rangle$ and a gradual semantics $\mathbf{S}$, an argument $a \in \mathcal{A}$ is threshold accepted in $\mathbf{G}$ with respect to $\mathbf{S}$ and some threshold $t \in [0,1]$, denoted $\mathbf{G} \mid\!\sim^t_{\mathbf{S}} a$, iff $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) \geq t$.*

Now that we can derive the acceptability of an argument from its strength, we may follow the same approach as Li, Oren, and Norman to determine the probability that a set of arguments is acceptable.

Definition 3.2 (Probability of Acceptability) *Given a probabilistic argumentation framework $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, a gradual semantics $\mathbf{S}$, and a threshold $t \in [0,1]$, the probability that some set of arguments $\mathcal{X} \subseteq \mathcal{A}$ is acceptable in $\mathbf{F}$ w.r.t. $\mathbf{S}$ and t is:*

$$P^{\mathbf{S},t}_{\mathbf{F}}(\mathcal{X}) = \sum_{\mathbf{G} \in \mathtt{I}(\mathbf{F}), \forall a \in \mathcal{X} \mathbf{G} \mid\!\sim^t_{\mathbf{S}} a} P^I_{\mathbf{F}}(\mathbf{G})$$

That is, $P^{\mathbf{S},t}_{\mathbf{F}}(\mathcal{X})$ is the sum of the probabilities of the induced graphs of $\mathbf{F}$ where all arguments in $\mathcal{X}$ are accepted. For brevity, we write $P^{\mathbf{S},t}_{\mathbf{F}}(x)$ instead of $P^{\mathbf{S},t}_{\mathbf{F}}(\mathcal{X})$ for singleton sets $\mathcal{X} = \{x\}$.

The following result states that the probability of acceptance of any argument in a framework is bounded from above by its probability of being present in the graph.

Proposition 3.3 *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics. For every $t \in [0,1]$ we have $P^{\mathbf{S},t}_{\mathbf{F}}(a) \leq P_{\mathcal{A}}(a)$.*

Now we state a form of monotonicity property that compares probabilities of acceptance when different thresholds are considered.

Proposition 3.4 *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics. If $t \leq t'$, then $P^{\mathbf{S},t}_{\mathbf{F}}(a) \geq P^{\mathbf{S},t'}_{\mathbf{F}}(a)$.*

Next we consider properties that depend on behaviour of the chosen gradual semantics. According to the first property, if the semantics satisfies Maximality, then the probability that a non-attacked argument has maximal strength is 1.

Proposition 3.5 *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics that satisfies Maximality. If $a \in \mathcal{A}$ is not attacked in $\mathbf{F}$, then $P^{\mathbf{S},1}_{\mathbf{F}}(a) = P_{\mathcal{A}}(a)$.*

On the other hand, Resilience implies non-zero probability that an argument has at least some positive strength.

Proposition 3.6 *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics that satisfies Resilience. Then for some $t > 0$, $P^{\mathbf{S},t}_{\mathbf{F}}(a) > 0$.*

If, in addition, the semantics satisfies Weakening, the probability of acceptance of an attacked argument cannot reach its probability of belonging to the graph.

Proposition 3.7 *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics that satisfies Resilience and Weakening. If a is attacked, then for $t = 1$, $P^{\mathbf{S},t}_{\mathbf{F}}(a) < P_{\mathcal{A}}(a)$.*

The last result of this section characterises perfect arguments.

Proposition 3.8 *Let $\mathbf{F}$ be a probabilistic argumentation framework and let $\mathbf{S}$ be a gradual semantics that satisfies Maximality, Weakening and Resilience. Then a is perfect in $\mathbf{F}$ iff $P^{\mathbf{S},1}_{\mathbf{F}}(a) = 1$.*

4 Ranking arguments in a probabilistic setting

Through the introduction, removal, or alteration of terms in a gradual semantics one may vastly alter the exact values it assigns to arguments. Hence, comparing an argument's strength to an exact threshold can only be informative when we are intimately familiar with the semantics used. We see this reflected in the principles used in the literature to study or define semantics [5, 11]; these principles only speak about strength values relative to those of other arguments, or to minimum or maximum values enforced by the framework and never do so in absolute values. Such properties are ultimately what distinguish two semantics and it is thus appropriate to compare arguments based on their strength.

Let us consider a practical example to better inform our intuitions. Suppose we are a university hiring committee tasked with filling a PhD position and there are two candidates to consider: Alex and Billy. We may model the suitability for hiring of Alex with an argument a and that of Billy with an argument b. Any reason for questioning the suitability of either candidate can now be modelled as an attacker of either argument. For instance, Alex's suitability may be brought into question based on doubts of their mastery of the English language (argument x with $x\mathcal{R}a$) and Billy may be considered a poor fit in the team as they are known to have had some conflict with other members of the research group (argument y with $y\mathcal{R}b$). Uncertainty is introduced into the graph by argument x (Alex has insufficient mastery of the English language) relying on an assumption made because Alex did not provide formal test results proving the contrary and by doubts whether argument y (there is known conflict between Billy and other team members) should constitute a valid attack on Billy's suitability as a candidate.

Say Alex is part of a demographic that is currently underrepresented in the university's staff, while Billy is not. Based on this fact, we prefer to hire Alex whenever they are at least as suitable a candidate as Billy. In order to select a candidate, we may wish to determine which candidate has the highest probability of being preferred. In other words, we want to find the probability that a is at least as strong as b and the probability that it is not and hire Alex if the former is greater.

Note that Definition 3.2 cannot be used to formalise this problem; while expressions of the form $P^{\mathbf{S},t}_{\mathbf{F}}(\mathcal{X})$ can be used to assign probability boundaries of arguments' acceptability, they do not allow for probability assignment of the explicit comparison of argument strengths. To find these probabilities, we first need to rank arguments in each individual induced framework. We now define what it means that a is ranked at least as highly as b in $\mathbf{G}$, denoted by $\mathbf{G} \vDash_{\mathbf{S}} a \preceq b$. Note that here we follow the convention from [2], where $a \preceq b$ means "a is at least as strong as b".

Definition 4.1 (Ranking Arguments in Induced Graphs) *Let $\mathbf{G}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ be an induced graph of probabilistic argumentation framework $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, and let $\mathbf{S}$ be a gradual semantics. For $a, b \in \mathcal{A}$, $\mathbf{G}' \vDash_{\mathbf{S}} a \preceq b$ iff one of the following holds:*

- *$a, b \in \mathcal{A}'$ and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(a) \geq \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(b)$, or*
- *either $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(b) = 0$ or $b \notin \mathcal{A}'$.*

The second condition equates the arguments without any strength with those not present in a graph. Note that now we can define probability of $a \preceq b$ as

$$P_{\mathbf{F},\mathbf{S}}(a \preceq b) = \sum_{\mathbf{G} \in \mathtt{I}(\mathbf{F}), \mathbf{G} \vDash_{\mathbf{S}} a \preceq b} P^I_{\mathbf{F}}(\mathbf{G}) \tag{2}$$

and that we get the probability of $a \not\preceq b$ as the complement of $a \preceq b$.

Now we emphasise that in some situations calculating probability that $a \preceq b$ is still insufficient. Suppose we introduce a third candidate into our example: Charlie, whose suitability is represented by argument c which also receives some uncertain attack. In this case the definition given by equation 2 no longer serves to determine the probability that, say, Alex is at least as strong candidate as the other two—even if we were to somehow combine $P_{\mathbf{F},\mathbf{S}}(a \preceq b)$ and $P_{\mathbf{F},\mathbf{S}}(a \preceq c)$—as induced graphs in which c is stronger than a may count toward $P_{\mathbf{F},\mathbf{S}}(a \preceq b)$ and graphs in which b is stronger than a may count toward $P_{\mathbf{F},\mathbf{S}}(a \preceq c)$. To properly determine this probability we require a formalism that considers multiple argument inequalities at the same time.

Definition 4.2 (Ranking Query) *Given a probabilistic argumentation framework $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, a ranking query is a Boolean combination of expressions of the form $a \preceq b$ with $a, b \in \mathcal{A}$.*

We will denote ranking queries with Greek letters $\alpha, \beta, \gamma, \dots$ Note that using ranking queries one can also express that one argument is of strictly higher rank than another ($a \preceq b \wedge \neg b \preceq a$) and that two arguments have equal ranks ($a \preceq b \wedge b \preceq a$). By $\top$ we denote a ranking query of the form $\alpha \vee \neg\alpha$.

For a ranking query α and an induced graph $\mathbf{G}$, we define $\mathbf{G} \vDash_{\mathbf{S}} \alpha$ simply by extending Definition 4.1 with the cases of Boolean connectives in the standard way. For example, we define $\mathbf{G} \vDash_{\mathbf{S}} \alpha \vee \beta$ iff $\mathbf{G} \vDash_{\mathbf{S}} \alpha$ or $\mathbf{G} \vDash_{\mathbf{S}} \beta$. We say that two queries α and β are incompatible if there is no graph such that $\mathbf{G} \vDash_{\mathbf{S}} \alpha$ and $\mathbf{G} \vDash_{\mathbf{S}} \beta$. Note that the query $\top$ holds in every induced graph.

Definition 4.3 (Probability of Ranking Queries) *Given a probabilistic argumentation framework $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and a gradual semantics $\mathbf{S}$, let α be a ranking query. The probability that the ranking on arguments indicated by α holds is:*

$$P_{\mathbf{F},\mathbf{S}}(\alpha) = \sum_{\mathbf{G} \in \mathrm{I}(\mathbf{F}), \mathbf{G} \vDash_{\mathbf{S}} \alpha} P^{I}_{\mathbf{F}}(\mathbf{G})$$

That is, $P_{\mathbf{F},\mathbf{S}}(x)$ is the sum of the probabilities of the induced graphs of $\mathbf{F}$ that entail the query α under semantics $\mathbf{S}$.

The additional expressivity offered by definition 4.3 allows us to successfully find the probability that Alex is at least as strong as the other candidates, namely by calculating $P_{\mathbf{F},\mathbf{S}}(a \preceq b \wedge a \preceq c)$.

The first part of the proposition below expresses finite additivity.

Proposition 4.4 *Given a PrAF $\mathbf{F}$ and a gradual semantics $\mathbf{S}$, let α be a ranking query.*

1. *If α and β are incompatible, then $P_{\mathbf{F},\mathbf{S}}(\alpha \vee \beta) = P_{\mathbf{F},\mathbf{S}}(\alpha) + P_{\mathbf{F},\mathbf{S}}(\beta)$.*
2. *$P_{\mathbf{F},\mathbf{S}}(\top) = 1$.*

The following result links Definition (4.3) and Definition (3.2).

Theorem 1 (Necessitation) *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ be a PrAF and $\mathbf{S}$ a gradual semantics. For all arguments $a, b \in \mathcal{A}$, if $P_{\mathbf{F},\mathbf{S}}(a \preceq b) = 1$, then for every $t \in [0, 1]$ we have $P^{\mathbf{S},t}_{\mathbf{F}}(a) \geq P^{\mathbf{S},t}_{\mathbf{F}}(b)$.*

5 A Constellations Approach to Argument Strength

In the previous sections we saw how we may use the probabilities of a PrAF's induced graphs together with the strength values assigned to each argument in those induced by a gradual semantics to determine the probability that an argument's strength meets some threshold or to determine the probability that some ranking query on arguments is satisfied. Considering each of the PrAF's induced graphs, as we did in answering both of these questions, is characteristic of the constellations approach to probabilistic argumentation. In this section, we explore the possibility of taking a similar approach in assigning each argument a unique strength; we look for a method that, given a choice of gradual semantics $\mathbf{S}$, assigns each argument an overall strength that considers both the probabilities of induced graphs and the strengths assigned by $\mathbf{S}$ in their context. First, we present a generalised notion of this new method for assigning strengths. Then, we discuss some of the properties and assumptions of this new approach and present a set of principles for them inspired by the principles for gradual semantics present in the literature. Finally, we propose a specification of the generalised method for assigning strengths based on a gradual semantics and discuss the properties of this specification.

Just as a gradual semantics is a function transforming an argumentation graph into a weighting on its elements, our generalised method for assigning unique strengths to arguments in a probabilistic argumentation framework—henceforth called a generalised semantics—is a function transforming a PrAF into a weighting on its arguments.

Definition 5.1 (Generalised Semantics) *A generalised semantics is a function $\mathbb{S}$ transforming any probabilistic argumentation framework $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle \in \mathtt{PrAF}$ into a weighting $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}$ on $\mathcal{A}$ (i.e., $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}} : \mathcal{A} \to [0, 1]$). For any $a \in \mathcal{A}$, $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a)$ represents the strength of a.*

5.1 Principles

As the PrAFs to which generalised semantics are applied are an extension of the AGs to which gradual semantics are applied, it is worth investigating which principles used in the study of gradual semantics (or at least the intuitions underlying them) transfer to the constellations approach. We study the generally desirable principles proposed in [5] which are recalled and adjusted to non-weighted graphs in Section 2.2 and find that while many translate naturally to the new setting, three require more extensive alteration to be sensible, and two are not generally desirable. Before proposing the principles resulting from this examination, let us consider these two principles, equivalence and reinforcement, that are not generally desirable in the new setting where we want both the probabilities of induced graphs and the strengths of arguments in them to contribute to the overall strength of an argument.

The intuition underlying the equivalence principle is that the strength of an argument in an argumentation graph should depend only on the strength of its direct attackers. The intuition underlying the reinforcement principle adds to this that increasing the strength of an attacker should increase the impact of its attack. To show how these intuitions may not generally hold when assuming the overall strength of attackers in a PrAF, we present the following: consider the PrAF $\mathbf{F}$ shown in Figure 1 with two induced graphs $\mathbf{G}$ and $\mathbf{G}'$ shown in Subfigures a and b respectively. Recall that we want the overall strength of an argument to be based on the strengths assigned to it in each induced graph by a gradual semantics $\mathbf{S}$. We are interested in the overall strengths of arguments a and b and how they are affected by their respective attackers x and y. If we we were to select h-Categoriser (def. 2.4) as $\mathbf{S}$, we would have $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a) = 1/2$, $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(a) = 1$, and $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(b) = \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(b) = 2/3$. The strength of y would be $1/2$ in all induced graphs and we would have $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(x) = 1$. Based on the

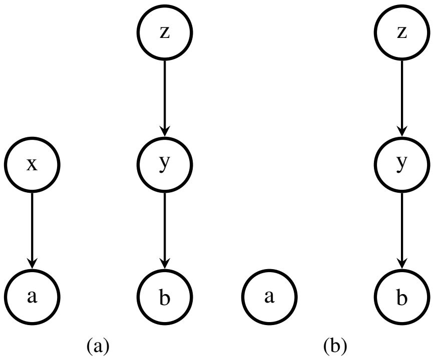

Figure 1: A probabilistic argumentation framework illustrating compensation between strength and probability where argument x is uncertain and all other elements are certain: (a) shows the entire PrAF and its induced graph where x is present; (b) shows the induced graph where x is not present.

neutrality principle's prescription that an argument with strength 0 contributes the same as no argument, we say $\mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}'}(x) = 0$.

Now that we know the strength of the arguments in each induced graph, the question becomes how to aggregate these. Given that, in the constellations approach, we consider each induced graph as a possible world with some probability, we might approach the matter similarly to an expected value and say that the amount an argument's strength in one induced graph contributes to its overall strength should be directly proportional to the probability of that induced graph. That is to say we multiply each strength in an induced graph with that graph's probability and sum over graphs to find the overall strength. If we take this approach, we may alter the value of $P_{\mathcal{A}}(x)$, and by extension $P^I_{\mathbf{F}}(\mathbf{G})$ and $P^I_{\mathbf{F}}(\mathbf{G}')$, to create scenarios where we may not desire equivalence or reinforcement.

First consider equivalence: say $P_{\mathcal{A}}(x) = 1/2$. This gives $P^I_{\mathbf{F}}(\mathbf{G}) = P^I_{\mathbf{F}}(\mathbf{G}') = 1/2$. The overall strength of x and y is equal, with $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(x) = 1 \cdot 1/2 + 0 \cdot 1/2 = 1/2$ and $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(y) = 1/2 \cdot 1/2 + 1/2 \cdot 1/2 = 1/2$. The attackers of a and b are thus equally strong overall, and the certainty with which they are attacked is also equal, but we have $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = 1/2 \cdot 1/2 + 1 \cdot 1/2 = 3/4$ while $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b) = 2/3 \cdot 1/2 + 2/3 \cdot 1/2 = 2/3$; a and b do not have the same overall strength, nor do we necessarily want them to.

To demonstrate how reinforcement is not always desirable we say $P_{\mathcal{A}}(x) = 2/3$. We now end up with $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(x) = 2/3 > \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(y) = 1/2$ and $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b) = 2/3$. The attacker of a is stronger overall than that of b, but it is reasonable to desire that a and b are equally strong overall.

Having seen which principles' intuitions do not transfer nicely to the constellations setting, we present seven principles based on those presented in [5]. For this, we require the notion of isomorphisms on PrAFs:

Definition 5.2 (PrAF Isomorphism) *Let $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and $\mathbf{F}' = \langle \mathcal{A}', P'_{\mathcal{A}}, \mathcal{R}', P'_{\mathcal{R}} \rangle$ be two PrAFs. An isomorphism from $\mathbf{F}$ to $\mathbf{F}'$ is a bijective function f from $\mathcal{A}$ to $\mathcal{A}'$ such that:*

- $\forall a \in \mathcal{A}, P_{\mathcal{A}}(a) = P'_{\mathcal{A}}(f(a))$;
- $\forall a, b \in \mathcal{A}, a\mathcal{R}b$ *iff* $f(a)\mathcal{R}'f(b)$;
- $\forall (a,b) \in \mathcal{R}, P_{\mathcal{R}}((a,b)) = P'_{\mathcal{R}}((f(a), f(b)))$

If $\mathbf{F} = \mathbf{F}'$, we call any isomorphism from $\mathbf{F}$ to $\mathbf{F}'$ an automorphism.

Principle 5.3 (PrAF Anonymity) *A generalised semantics $\mathbb{S}$ satisfies anonymity iff for any two PrAFs $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and $\mathbf{F}' = \langle \mathcal{A}', P'_{\mathcal{A}}, \mathcal{R}', P'_{\mathcal{R}} \rangle$, and any isomorphism f from $\mathbf{F}$ to $\mathbf{F}'$, the following holds: $\forall a \in \mathcal{A}, \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}'}(f(a))$.*

Principle 5.4 (PrAF Independence) *A generalised semantics $\mathbb{S}$ satisfies independence iff for any two PrAFs $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and $\mathbf{F}' = \langle \mathcal{A}', P'_{\mathcal{A}}, \mathcal{R}', P'_{\mathcal{R}} \rangle$ where $\mathcal{A} \cap \mathcal{A}' = \emptyset$, it holds that: $\forall a \in \mathcal{A}, \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F} \oplus \mathbf{F}'}(a)$.*

Principle 5.5 (PrAF Directionality) *A generalised semantics $\mathbb{S}$ satisfies directionality iff for any two PrAFs $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and $\mathbf{F}' = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}', P'_{\mathcal{R}} \rangle$ where $\mathcal{R}' = \mathcal{R} \cup \{(a,b)\}$ and $\forall r \in \mathcal{R}, P'_{\mathcal{R}}(r) = P_{\mathcal{R}}(r)$, the following holds: for any $x \in \mathcal{A}$, if there is no path from b to x, $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(x) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}'}(x)$.*

Principle 5.6 (PrAF Maximality) *A generalised semantics $\mathbb{S}$ satisfies probability maximality iff for any PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a \in \mathcal{A}$, it holds that: if $\mathtt{Att}_{\mathbf{F}}(a) = \emptyset$, then $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = P_{\mathcal{A}}(a)$.*

Principle 5.7 (PrAF Weakening) *A generalised semantics $\mathbb{S}$ satisfies weakening iff for any PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a \in \mathcal{A}$, it holds that: if $\exists b \in \mathtt{Att}_{\mathbf{F}}(a)$ s.t. $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b) > 0$ then $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) < P_{\mathcal{A}}(a)$.*

Principle 5.8 (PrAF Weakening Soundness) *A generalised semantics $\mathbb{S}$ satisfies weakening soundness iff for any PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a \in \mathcal{A}$, the following holds: if $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) < P_{\mathcal{A}}(a)$ then $\exists b \in \mathtt{Att}_{\mathbf{F}}(a)$ s.t. $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b) > 0$.*

Principle 5.9 (PrAF Resilience) *A generalised semantics $\mathbb{S}$ satisfies resilience iff for any PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a \in \mathcal{A}$, $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) > 0$.*

We have already discussed that if we follow the constellation approach, we should not expect that the Equivalence principle holds. Now we turn to a special case of Equivalence proposed in [8], called Symmetry. Originally it states that two arguments a and b that have the same sets of attackers also have equal strengths. Similarly as for Equivalence, in the case of PrAFs that "symmetry" between a and b breaks easily when we zoom in to induced graphs. For example, if a and b are not certain and there are different (asymmetric) attacks from a and b towards their own attackers, the induced graph in which only a appears will be very different from those in which b occurs. In order to enforce symmetry we require that a and b interact with their attackers, attackers of their attackers, and so on in a symmetric way.

Definition 5.10 (Attack Structure) *For given PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$ and $a \in \mathcal{A}$, the attack structure of a in $\mathbf{F}$ is $\mathtt{Str}_{\mathbf{F}}(a) = \{a\} \cup \{c \in \mathcal{A} \mid$ there is a path from c to $a\}$. We denote by $\mathtt{Str}_{\mathbf{F}}(a,b)$ the set $\mathtt{Str}_{\mathbf{F}}(a) \cup \mathtt{Str}_{\mathbf{F}}(b)$.*

Principle 5.11 (PrAF Symmetry) *A generalised semantics $\mathbb{S}$ satisfies PrAF Symmetry iff for every PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if $f : \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)} \to \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)}$ s.t. $f(a) = b$, $f(b) = a$ and $f(x) = x$ otherwise, is an automorphism, then $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b)$.*

The principles Neutrality and Counting build on the idea of symmetry, and say that if an attack is added to one of two arguments in a symmetric situation, the attack will additionally harm the target if the strength of that attacker is positive, otherwise it will not. We apply that intuition directly to our notion of PrAF symmetry.

Principle 5.12 (PrAF Neutrality) *A generalised semantics $\mathbb{S}$ satisfies Praf Neutrality iff for every PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if there exists c such that $(c, b) \in \mathcal{R}$, $(c, b) \notin \mathcal{R}$, and $f : \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)\setminus\{c\}} \to \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)\setminus\{c\}}$ s.t. $f(a) = b$, $f(b) = a$ and $f(x) = x$ otherwise, is an automorphism, then $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(c) = 0$ implies $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b)$.*

Principle 5.13 (PrAF Counting) *A generalised semantics $\mathbb{S}$ satisfies PrAF Counting iff for every PrAF $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, $\forall a, b \in \mathcal{A}$, the following holds: if there exists c such that $(c, b) \in \mathcal{R}$, $(c, b) \notin \mathcal{R}$, and $f : \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)\setminus\{c\}} \to \mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a,b)\setminus\{c\}}$ s.t. $f(a) = b$, $f(b) = a$ and $f(x) = x$ otherwise, is an automorphism, then $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(c) > 0$ implies $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) > \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(b)$.*

The following result shows that PrAF Symmetry is already a consequence of a subset of other principles.

Theorem 2 *If a generalised semantics $\mathbb{S}$ satisfies PrAF Anonymity, PrAF Independence and PrAF Directionality, then $\mathbb{S}$ also satisfies PrAF Symmetry.*

We now present the result which claims that from a subset of the principles it follows that an argument's strength is bounded by its probability.

Theorem 3 *If a generalised semantics $\mathbb{S}$ satisfies PrAF Independence, PrAF Maximality, PrAF Weakening, V Neutrality, and PrAF Directionality, then for any $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle \in \mathtt{PrAF}$, for any $a \in \mathcal{A}$, $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) \leq P_{\mathcal{A}}(a)$.*

Our next formal result states that there is always a subset of arguments in a probabilistic argumentation framework which impacts the strength of a given argument, namely its attack structure.

Theorem 4 *If a semantics $\mathbb{S}$ satisfies PrAF-Independence and PrAF-Directionality, then for any $\mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle$, for any $a \in \mathcal{A}$, the following holds: $\mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}}(a) = \mathtt{Deg}^{\mathbb{S}}_{\mathbf{F}_{|\mathtt{Str}_{\mathbf{F}}(a)}}(a)$.*

5.2 *Expected strength semantics*

With a set of principles—or desirable properties of a generalised semantics—laid out, let us consider one way of specifying a generalised semantics $\mathbb{S}$ using a gradual semantics $\mathbf{S}$ and consider how this specification relates to the different principles. When we were discussing equivalence and reinforcement, we weighted the contribution of an argument's strength in an induced graph with that induced graph's probability. Formalising this gives:

Definition 5.14 (Expected Strength Semantics) *Given a gradual semantics $\mathbf{S}$, the expected strength semantics based on $\mathbf{S}$, denoted $E(\mathbf{S})$, is the generalised semantics such that $\forall \mathbf{F} = \langle \mathcal{A}, P_{\mathcal{A}}, \mathcal{R}, P_{\mathcal{R}} \rangle \in \mathtt{PrAF}, \forall a \in \mathcal{A}$,*

$$\mathtt{Deg}^{E(\mathbf{S})}_{\mathbf{F}}(a) = \sum_{\mathbf{G}=\langle \mathcal{A}', \mathcal{R}' \rangle \in \mathtt{I}(\mathbf{F}), a \in \mathcal{A}'} P^{I}_{\mathbf{F}}(\mathbf{G}) \cdot \mathtt{Deg}^{\mathbf{S}}_{\mathbf{G}}(a)$$

The next results show that if a semantics $\mathbf{S}$ satisfies some principles from the literature (presented in Section 2.2), then $E(\mathbf{S})$ satisfies the principles for generalised semantics proposed in this section.

Theorem 5 *Let $\mathbf{S}$ be a gradual semantics.*

- *If $\mathbf{S}$ satisfies Anonymity, then $E(\mathbf{S})$ satisfies PrAF Anonymity.*
- *If $\mathbf{S}$ satisfies Independence, then $E(\mathbf{S})$ satisfies PrAF Independence.*
- *If $\mathbf{S}$ satisfies Directionality, then $E(\mathbf{S})$ satisfies PrAF Directionality.*
- *If $\mathbf{S}$ satisfies Maximality, then $E(\mathbf{S})$ satisfies PrAF Maximality.*
- *If $\mathbf{S}$ satisfies Weakening and Resilience, then $E(\mathbf{S})$ satisfies PrAF Weakening.*
- *If $\mathbf{S}$ satisfies Weakening Soundness, then $E(\mathbf{S})$ satisfies PrAF Weakening Soundness.*
- *If $\mathbf{S}$ satisfies Resilience, then $E(\mathbf{S})$ satisfies PrAF Resilience.*
- *If $\mathbf{S}$ satisfies Anonymity, Independence and Directionality, then $E(\mathbf{S})$ satisfies PrAF Symmetry.*
- *If $\mathbf{S}$ satisfies Anonymity, Independence, Directionality and Neutrality, then $E(\mathbf{S})$ satisfies PrAF Neutrality.*
- *If $\mathbf{S}$ satisfies Anonymity, Independence, Directionality, Counting and Resilience, then $E(\mathbf{S})$ satisfies PrAF Counting.*

Let us recall that the h-categoriser semantics satisfies all the principles from Section 2.2 [5]. Together with Theorem 5, that gives us the following result, which verifies compatibility of our principles.

Theorem 6 $E(\mathtt{Hbs})$ *satisfies all the principles proposed in Sec. 5.1.*

6 Conclusion

In this article, we conducted the first study of gradual semantics in probabilistic argumentation frameworks. Following the constellations approach, we defined the probability of an argument's acceptability with respect to an arbitrary threshold that correspond to possible strength of the argument, and probabilities of rankings of arguments. We also proposed a method to calculate the overall strength of each argument in the framework, and we evaluated this method against a set of principles. Following the original approach from [28], we assumed probabilistic independence of elements of the framework belonging to a graph. It is easy to modify our approach to the more general case, where the independency assumptions are not used (like, for example, in [23]). Instead of calculating the probability of an induced graph from the probabilities of the arguments and attacks (Definition 2.6), that probability has to be taken as a constituent of the framework. Essentially, nothing would change in our approach: instead of using Definition 2.6 to calculate probabilities of acceptability ($P^{\mathbf{S},t}_{\mathbf{F}}(\mathcal{X})$) or a ranking query holding ($P_{\mathbf{F},\mathbf{S}}(\alpha)$), we would simply use probabilities of induced graphs given as a parts of the framework.

It is worth mentioning the work of Thimm, Cerutti and Rienstra [38] on gradual semantics based on constellations approach to probabilistic argumentation. That work is significantly different than ours: they did not study the application of gradual semantics to PrAFs. Instead, they proposed a gradual semantics for standard argumentation frameworks, using PrAFs and existing Dung semantics as a tool.

When proposing the set of principles, we generalised existing existing principles from [5], which refine the first set of principles for gradual semantics [3], and extend them to the weighted argumentation frameworks. Here we considered flat (non-weighed) frameworks, so we could have based our generalised principles on [3] as well. We didn't find an advantage of using [5] instead of [3], since two sets of principles describe same intuitions. On the other hand, our intention is to investigate generalisation of our work to weighted graphs as future work, in which case we can directly retrieve original variants of principles from [5]. We believe that probabilities and weights can coexist in a meaningful way, as the intuitions behind them are coherent: where probabilities encode the likelihood of an element appearing in a graph, weights can encode the basic strength of an argument.

References

[1] Leila Amgoud, ‘A Replication Study of Semantics in Argumentation.’, in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pp. 6260–6266, (2019).

[2] Leila Amgoud and Jonathan Ben-Naim, ‘Ranking-Based Semantics for Argumentation Frameworks.’, in *Scalable Uncertainty Management - 7th International Conference, SUM 2013, Washington, DC, USA, September 16-18, 2013. Proceedings*, pp. 134–147, (2013).

[3] Leila Amgoud and Jonathan Ben-Naim, ‘Axiomatic Foundations of Acceptability Semantics.’, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016.*, pp. 2–11, (2016).

[4] Leila Amgoud and Jonathan Ben-Naim, ‘Evaluation of arguments in weighted bipolar graphs.’, *Int. J. Approx. Reason.*, **99**, 39–55, (2018).

[5] Leila Amgoud, Jonathan Ben-Naim, Dragan Doder, and Srdjan Vesic, ‘Acceptability Semantics for Weighted Argumentation Frameworks.’, in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pp. 56–62, (2017).

[6] Leila Amgoud and Dragan Doder, ‘Gradual Semantics for Weighted Graphs: An Unifying Approach.’, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018.*, pp. 613–614, (2018).

[7] Leila Amgoud and Dragan Doder, ‘Gradual Semantics Accounting for Varied-Strength Attacks.’, in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’19, Montreal, QC, Canada, May 13-17, 2019*, pp. 1270–1278, (2019).

[8] Leila Amgoud, Dragan Doder, and Srdjan Vesic, ‘Evaluation of argument strength in attack graphs: Foundations and semantics.’, *Artif. Intell.*, **302**, 103607, (2022).

[9] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Ricardo Simari, Matthias Thimm, and Serena Villata, ‘Towards Artificial Argumentation.’, *AI Mag.*, **38**(3), 25–36, (2017).

[10] *Handbook of formal argumentation*, eds., Pietro Baroni, Dov M. Gabbay, Massimiliano Giacomin, and Leendert van der Torre, College Publications, 2018.

[11] Pietro Baroni, Antonio Rago, and Francesca Toni, ‘From fine-grained properties to broad principles for gradual argumentation: A principled spectrum.’, *Int. J. Approx. Reason.*, **105**, 252–286, (2019).

[12] Trevor J. M. Bench-Capon and Paul E. Dunne, ‘Argumentation in artificial intelligence.’, *Artif. Intell.*, **171**(10-15), 619–641, (2007).

[13] Philippe Besnard and Anthony Hunter, ‘A logic-based theory of deductive arguments.’, *Artif. Intell.*, **128**(1-2), 203–235, (2001).

[14] Stefano Bistarelli, Theofrastos Mantadelis, Francesco Santini, and Carlo Taticchi, ‘Probabilistic Argumentation Frameworks with MetaProbLog and ConArg.’, in *IEEE 30th International Conference on Tools with Artificial Intelligence, ICTAI 2018, 5-7 November 2018, Volos, Greece.*, pp. 675–679, (2018).

[15] Stefano Bistarelli and Francesco Santini, ‘A Definition of Sceptical Semantics in the Constellations Approach.’, in *Logic Programming and Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022, Genova, Italy, September 5-9, 2022, Proceedings*, pp. 62–74, (2022).

[16] Elise Bonzon, Jérôme Delobelle, Sébastien Konieczny, and Nicolas Maudet, ‘A Comparative Study of Ranking-Based Semantics for Abstract Argumentation.’, in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pp. 914–920, (2016).

[17] Claudette Cayrol and Marie-Christine Lagasquie-Schiex, ‘Graduality in Argumentation.’, *J. Artif. Intell. Res.*, **23**, 245–297, (2005).

[18] Phan Minh Dung, ‘On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games.’, *Artif. Intell.*, **77**(2), 321–358, (1995).

[19] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, ‘Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence.’, *Artif. Intell.*, **268**, 1–29, (2019).

[20] Dov M. Gabbay and Odinaldo Rodrigues, ‘Equilibrium States in Numerical Argumentation Networks.’, *Logica Universalis*, **9**(4), 411–473, (2015).

[21] Anthony Hunter, ‘Some Foundations for Probabilistic Abstract Argumentation.’, in *Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012*, pp. 117–128, (2012).

[22] Anthony Hunter, ‘A probabilistic approach to modelling uncertain logical arguments.’, *Int. J. Approx. Reason.*, **54**(1), 47–81, (2013).

[23] Anthony Hunter, ‘Probabilistic qualification of attack in abstract argumentation.’, *Int. J. Approx. Reason.*, **55**(2), 607–638, (2014).

[24] Anthony Hunter and Kawsar Noor, ‘Aggregation of Perspectives Using the Constellations Approach to Probabilistic Argumentation.’, in *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020.*, pp. 2846–2853, (2020).

[25] Anthony Hunter and Matthias Thimm, ‘Optimization of dialectical outcomes in dialogical argumentation.’, *Int. J. Approx. Reason.*, **78**, 73–102, (2016).

[26] Anthony Hunter and Matthias Thimm, ‘Probabilistic Reasoning with Abstract Argumentation Frameworks.’, *J. Artif. Intell. Res.*, **59**, 565–611, (2017).

[27] João Leite and João G. Martins, ‘Social Abstract Argumentation.’, in *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pp. 2287–2292, (2011).

[28] Hengfei Li, Nir Oren, and Timothy J. Norman, ‘Probabilistic Argumentation Frameworks.’, in *Theorie and Applications of Formal Argumentation - First International Workshop, TAFA 2011. Barcelona, Spain, July 16-17, 2011, Revised Selected Papers*, pp. 1–16, (2011).

[29] Fangzhen Lin and Yoav Shoham, ‘Argument Systems: A Uniform Basis for Nonmonotonic Reasoning.’, in *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (KR’89). Toronto, Canada, May 15-18 1989.*, pp. 245–255, (1989).

[30] Theofrastos Mantadelis and Stefano Bistarelli, ‘Probabilistic abstract argumentation frameworks, a possible world view.’, *Int. J. Approx. Reason.*, **119**, 204–219, (2020).

[31] Kawsar Noor and Anthony Hunter, ‘Analysing Product Reviews Using Probabilistic Argumentation.’, in *Computational Models of Argument - Proceedings of COMMA 2020, Perugia, Italy, September 4-11, 2020.*, pp. 295–306, (2020).

[32] Célia da Costa Pereira, Andrea Tettamanzi, and Serena Villata, ‘Changing One’s Mind: Erase or Rewind?’, in *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pp. 164–171, (2011).

[33] Sylwia Polberg and Anthony Hunter, ‘Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches.’, *Int. J. Approx. Reason.*, **93**, 487–543, (2018).

[34] Nico Potyka, ‘Extending Modular Semantics for Bipolar Weighted Argumentation.’, in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’19, Montreal, QC, Canada, May 13-17, 2019*, pp. 1722–1730, (2019).

[35] Nico Potyka, ‘Attractor - A Java Library for Gradual Bipolar Argumentation.’, in *Computational Models of Argument - Proceedings of COMMA 2022, Cardiff, Wales, UK, 14-16 September 2022*, pp. 369–370, (2022).

[36] Christian Straßer and Lisa Michajlova, ‘Evaluating and selecting arguments in the context of higher order uncertainty’, *Frontiers in Artificial Intelligence*, **6**, (2023).

[37] Matthias Thimm, ‘A Probabilistic Semantics for abstract Argumentation.’, in *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*, pp. 750–755, (2012).

[38] Matthias Thimm, Federico Cerutti, and Tjitze Rienstra, ‘Probabilistic Graded Semantics.’, in *Computational Models of Argument - Proceedings of COMMA 2018, Warsaw, Poland, 12-14 September 2018*, pp. 369–380, (2018).

[39] Bruno Yun and Srdjan Vesic, ‘Gradual Semantics for Weighted Bipolar SETAFs.’, in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 16th European Conference, ECSQARU 2021, Prague, Czech Republic, September 21-24, 2021, Proceedings*, pp. 201–214, (2021).

[40] Qiaoting Zhong, Xiuyi Fan, Xudong Luo, and Francesca Toni, ‘An explainable multi-attribute decision model based on argumentation.’, *Expert Syst. Appl.*, **117**, 42–61, (2019).