Typicality, Conditionals and a Probabilistic Semantics for Gradual Argumentation

Mario Alviano[1,*], Laura Giordano[2,*] and Daniele Theseider Dupré[2,*]

[1]DEMACS, University of Calabria, Via Bucci 30/B, 87036 Rende (CS), Italy
[2]DISIT, University of Piemonte Orientale, Viale Michel 11, 15121 Alessandria, Italy

Abstract
In this paper we propose a general approach to define a many-valued preferential interpretation of gradual argumentation semantics. The approach allows for conditional reasoning over arguments and boolean combination of arguments, with respect to some chosen gradual semantics, through the verification of graded (strict or defeasible) implications over a preferential interpretation. The paper also develops and discusses a probabilistic semantics for gradual argumentation, which builds on the many-valued conditional semantics.

1. Introduction

Argumentation is one of the major fields in non-monotonic reasoning (NMR) which has been shown to be very relevant for decision making and for explanation [1]. The relationships between preferential semantics of commonsense reasoning [2, 3, 4, 5] and argumentation semantics are very strong [6, 4]. While for Dung-style argumentation semantics and for Abstract Dialectical Frameworks, the relationships with conditional reasoning have been deeply investigated [7, 8, 9, 10], this is not the case for gradual argumentation [11, 12, 13, 14, 15, 16, 17].

In a companion paper [18], we have proposed an ASP approach for conditional reasoning over weighted argumentation graphs in a specific gradual semantics (the φ-coherent semantics), through the verification of graded conditional implications over arguments and over boolean combinations of arguments. In this paper, we show that the proposal can be generalized to a larger class of gradual argumentation semantics.

The paper proposes a general approach to define a preferential interpretation of an argumentation graph under a gradual semantics (provided weak conditions on the domain of argument interpretation are satisfied), to allow for *conditional reasoning over the argumentation graph*, by formalizing conditional properties of the graph in a many-valued logic with typicality: a many-valued propositional logic in which arguments play the role of propositional variables, and a typicality operator is allowed, inspired by the typicality operator proposed in the Propositional Typicality Logic [19] as well as in Description Logics (DLs) with typicality [20]. The operator allows for the definition of *conditional implications* $\mathbf{T}(A_1) \rightarrow A_2$, meaning that "normally argument A_1 implies argument A_2", in the sense that "in the typical situations where A_1 holds, A_2 also holds". The truth degree of such implications can be determined with respect to a preferential interpretation defined from a set of labellings of an argumentation graph, according to a chosen (gradual) argumentation semantics. They correspond to conditional implications $\alpha \mathrel{|\!\sim} \beta$ in the KLM approach [21, 3]. More precisely, in this paper we consider *graded conditionals* of the form $\mathbf{T}(\alpha) \rightarrow \beta \geq l$, meaning that "normally argument α implies argument β with degree at least l", where α and β can be *boolean combination of arguments*. They are inspired by graded inclusion axioms in fuzzy DLs [22] and in weighted defeasible knowledge bases in DLs [23]. The satisfiability of such implications in the *multi-preferential* interpretation I^S_G of an argumentation graph G (wrt. a given semantics S), exploits multiple preference relations $<_{A_i}$ over labellings, each one associated with an argument A_i.

We reformulate the KLM postulates of a preferential consequence relation for graded conditionals and prove that they are satisfied by the conditionals which hold in the multi-preferential interpretation I^S_G, for some choice of combination functions. We also prove that the satisfiability of a graded conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ in a *finite* preferential interpretation I^S_G can be decided in polynomial time in the size of the interpretation I^S_G times the size of the conditional formula.

The definition of a preferential interpretation I^S_G associated with an argumentation graph G and a gradual semantics S also sets the ground for the definition of a *probabilistic interpretation* of gradual semantics S. For the gradual semantics with domain of argument valuation

21st International Workshop on Nonmonotonic Reasoning, September 2–4, 2023, Rhodes, Greece
*Corresponding author.
mario.alviano@unical.it (M. Alviano); laura.giordano@uniupo.it (L. Giordano); dtd@uniupo.it (D. Theseider Dupré)
https://alviano.net/ (M. Alviano); https://people.unipmn.it/laura.giordano/ (L. Giordano); https://people.unipmn.it/dtd/ (D. Theseider Dupré)
0000-0002-2052-2063 (M. Alviano); 0000-0001-9445-7770 (L. Giordano); 0000-0001-6798-4380 (D. Theseider Dupré)

in the unit real interval $[0, 1]$, we propose a probabilistic argumentation semantics, which builds on a gradual semantics and is inspired by Zadeh's probability of fuzzy events [24]. As we will see it can be regarded as a generalization of the probabilistic semantics in [25] to the gradual case.

2. A preferential interpretation of gradual argumentation semantics

Given an argumentation graph G and some gradual argumentation semantics S, we define a preferential (many-valued) interpretation of the argumentation graph G, with respect to the gradual semantics S. We generalize the approach proposed in [18] for weighted argumentation graphs, without assuming a specific gradual semantics. In the following, we will consider both weighted and non-weighted argumentation graphs.

We follow Baroni, Rago and Toni [16, 26] (in their definition of a Quantitative Bipolar Argumentation Framework, QBAF) in the choice of the *domain of argument interpretation*, letting it to be a set $\mathcal{D}$, equipped with a *preorder relation* $\leq$, an assumption which is considered general enough to include the domain of argument valuations in most gradual argumentation semantics. As usual, we let $x < y$ iff $x \leq y$ and $y \not\leq x$.

As in [16], we do not assume $\mathcal{D}$ contains a *minimum element* and a *maximum element*. However, if a *minimum element* and a *maximum element* belong to $\mathcal{D}$, we will denote them by $0_\mathcal{D}$ and $1_\mathcal{D}$ (or simply 0 and 1), respectively. If not, we will add the two elements $0_\mathcal{D}$ and $1_\mathcal{D}$ at the bottom and top of the values in $\mathcal{D}$, respectively. We will also call $\mathcal{D}$ the *truth value set* (or the *truth degree set*). For instance, $\mathcal{D}$ may be unit interval $[0, 1]$ or, in the finitely-valued case (as in [18]), the finite set $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, 1\}$, for some integer $n \geq 1$.

For the definition of an argumentation graph, we consider the definition of *edge-weighted QBAF* by [27], for a generic domain $\mathcal{D}$. As we want to capture both weighted and non-weighted argumentation graphs, in the following, we will let the label of edges of the graph be $+1$ or -1 to denote support and attack in the non-weighted case (see below).

We let a *(weighted) argumentation graph* to be a quadruple $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$, where $\mathcal{A}$ is a set of arguments, $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ a set of edges, $\sigma_0 : \mathcal{A} \rightarrow \mathcal{D}$ assigns a *base score* of arguments, and $\pi : \mathcal{R} \rightarrow \mathbb{R}$ is a weight function assigning a positive or negative weight to edges. An example of weighted argumentation graph is in Figure 1, where the base score (i.e., the initial valuation of arguments) is not represented.

A pair $(B, A) \in \mathcal{R}$ is regarded as a *support* of argument B to argument A when the weight $\pi(B, A)$ is positive and as an *attack* of argument B to argument A when $\pi(B, A)$ is negative. In case the graph is *non-weighted*, we let $\pi(B, A) = -1$ mean that argument B attacks argument A, and $\pi(B, A) = +1$ mean that argument B supports argument A.

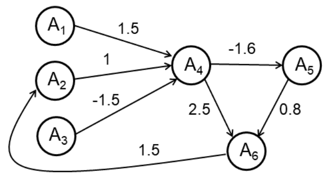

Figure 1: Example weighted argumentation graph G from [18]

Bipolar argumentation has been studied in the literature [28, 16, 26, 27] through different frameworks. We refer to the Quantitative Bipolar Argumentation Framework (QBAF) by Baroni, Rago and Toni [16, 26] for a classification and the properties of gradual semantics, when the argumentation graph is non-weighted, and to Potyka's work [27] for the framework of edge-weighted QBAFs and its properties. The properties of edge-weighted argumentation graphs with weights in the unit interval $[0, 1]$ have as well been studied in the gradual semantics framework by Amgoud and Doder [17].

Whatever semantics S is considered for an argumentation graph G, we will assume that S identifies *a set* Σ^S *of labellings* of the graph G over a domain of argument valuation $\mathcal{D}$. A *labelling* σ *of* G *over* $\mathcal{D}$ is a function $\sigma : \mathcal{A} \rightarrow \mathcal{D}$, which assigns to each argument an *acceptability degree* (or a *strength*) in the domain of argument valuation $\mathcal{D}$.[1] In some cases, we may omit the base score σ_0, and consider the set of labellings Σ^S of a graph $G = \langle \mathcal{A}, \mathcal{R}, \pi \rangle$, for all the possible choices of the base score, or a subset of them.

As an example, we refer to (without providing its definition) the φ-coherent semantics [29, 30] of graph G in Figure 1.

Example 1 ([18]). *As an example, in the φ-coherent semantics for weighted argumentation graphs, in the finitely-valued case, for $\mathcal{D} = \mathcal{C}_n$ with $n = 5$, the graph G in Figure 1 has 36 labellings, while, for $n = 9$, G has 100 labellings. For instance, $\sigma = (0, 4/5, 3/5, 2/5, 2/5, 3/5)$ (meaning that $\sigma(A_1) = 0$, $\sigma(A_2) = 4/5$, and so on) is a*

[1] Clearly, not any mapping qualifies as a labelling of a gradual semantics S, as a gradual semantics is intended to satisfy some principles, such as those identified in the different frameworks mentioned above [16, 26, 17, 27] for the non-weighted and for the weighted case. For our concerns, in the following, we will assume that, whatever the concrete definition of a semantics S might be, the semantics S can be regarded, abstractly, as a pair $(\mathcal{D}, \Sigma^S)$: a domain of argument valuation $\mathcal{D}$ and a set of labellings Σ^S over the domain.

labelling for $n = 5$.

In the following, we introduce a propositional language to represent *boolean combination of arguments* and a *many-valued semantics* for it over the domain $\mathcal{D}$ of argument valuation. Then, we extend the language with a *typicality operator*, to introduce defeasible implications over boolean combinations of arguments and define a *(multi-)preferential* interpretation associated with the argumentation graph G and a set of labellings Σ^S.

Given an argumentation graph $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$, we introduce a propositional language $\mathcal{L}$, whose set of propositional variables $Prop$ is the set of arguments $\mathcal{A}$. We assume our language $\mathcal{L}$ contains the connectives $\wedge$, $\vee$, $\neg$ and $\rightarrow$, and that formulas are defined inductively, as usual. Formulas built from the propositional variables in $\mathcal{A}$ correspond to a *boolean combination of arguments* (denoted α, β, γ), which are considered, for instance, by Hunter, Polberg and Thimm in [31].

We consider a many-valued semantics for boolean combination of arguments, with $\mathcal{D}$ as the *truth degree set*. Let $\otimes$, $\oplus$, $\rhd$ and $\ominus$ be the *truth degree functions* in $\mathcal{D}$ for the connectives $\wedge$, $\vee$, $\neg$ and $\rightarrow$ (respectively). When $\mathcal{D}$ is $[0, 1]$ or the finite set $\mathcal{C}_n$, as in our case of study [18], $\otimes$, $\oplus$, $\rhd$ and $\ominus$ can be chosen as a t-norm, s-norm, implication function, and negation function in some system of many-valued logic [32]; for instance, in Gödel logic, that we will consider later, $a \otimes b = min\{a, b\}$, $a \oplus b = max\{a, b\}$, $a \rhd b = 1$ *if* $a \leq b$ *and* b *otherwise*; and $\ominus a = 1$ *if* $a = 0$ *and* 1 *otherwise*.

A labelling $\sigma : \mathcal{A} \rightarrow \mathcal{D}$ of graph G, assigning to each argument $A_i \in \mathcal{A}$ a truth degree in $\mathcal{D}$, can be regarded as a *many-valued valuation*. A valuation σ can be inductively extended to all propositional formulas of $\mathcal{L}$ as follows:

$\sigma(\alpha \wedge \beta) = \sigma(\alpha) \otimes \sigma(\beta)$ $\quad$ $\sigma(\alpha \vee \beta) = \sigma(\alpha) \oplus \sigma(\beta)$
$\sigma(\alpha \rightarrow \beta) = \sigma(\alpha) \rhd \sigma(\beta)$ $\quad$ $\sigma(\neg \alpha) = \ominus \sigma(\alpha)$

Based on the choice of the combination functions, a labelling σ uniquely assigns a truth degree to any boolean combination of arguments. We will assume that the false argument $\bot$ and the true argument $\top$ are formulas of L and that $\sigma(\bot) = 0_{\mathcal{D}}$ and $\sigma(\top) = 1_{\mathcal{D}}$, for all labellings σ.

Definition 1. *Given a set of labellings Σ, for each argument $A_i \in \mathcal{A}$, we define a* preference relation $<_{A_i}$ *on Σ, as follows: for $\sigma, \sigma' \in \Sigma$,*

$$\sigma <^{\Sigma}_{A_i} \sigma' \textit{ iff } \sigma'(A_i) < \sigma(A_i),$$

Labelling σ *is preferred to* σ' with respect to argument A_i (or σ is *more plausible* than σ' for argument A_i), when the degree of truth of A_i in σ is greater than the degree of truth of A_i in σ'. The preference relation $<^{\Sigma}_{A_i}$ is a *strict partial order relation* on Σ. We will be simply write $<_{A_i}$, omitting Σ, when it is clear from the context.

The definition of preference over arguments which is induced by a set of labellings Σ also extends to the *boolean combination of arguments* α in the obvious way, based on the choice of combination functions. A set of labellings Σ *induces* a preference relation $<_\alpha$ on Σ, for each boolean combination of arguments α, as follows: for all $\sigma, \sigma' \in \Sigma$,

$$\sigma <_\alpha \sigma' \text{ iff } \sigma'(\alpha) < \sigma(\alpha).$$

When the set Σ^S of labellings of a graph in an argumentation semantics S is infinite, the preference relations $<_{A_i}$ (and $<_\alpha$) are not guaranteed to be *well-founded*, as there may be infinitely-descending chains of labellings. In the following, we will restrict our consideration to set of labellings Σ^S such that both $<_{A_i}$ and $<_{\neg A_i}$ are well-founded for all arguments A_i in Σ. We will call such a set of labellings Σ^S a *well-founded set of labellings*. From the monotonicity properties of t-norms and s-norms, and the antitonicity property of negation functions, it follows that, for any well-founded set of labellings Σ^S, $<_\alpha$ is also well-founded for any boolean combinations of arguments α.

We can now define the preferential interpretation of a graph with respect to a set of labellings.

Definition 2. *Given an argumentation graph G, a gradual semantics S with domain of argument valuation $\mathcal{D}$, and the set of labellings Σ^S of G wrt S, we let the* preferential interpretation of G wrt S *to be the pair* $I^S_G = (\mathcal{D}, \Sigma^S)$.

The preference relations $<_\alpha$ in the preferential interpretation I^S_G are left implicit, as they are induced by the labellings in Σ^S (according to Definition 1). Often, we will simply write I^S or I, rather than I^S_G.

Language $\mathcal{L}^{\mathbf{T}}$ is obtained by extending language $\mathcal{L}$ with a unary typicality operator $\mathbf{T}$. Intuitively, "a sentence of the form $\mathbf{T}(\alpha)$ is understood to refer to the *typical situations in which* α *holds*" [19]. The typicality operator allows the formulation of *conditional implications* (or *defeasible implications*) of the form $\mathbf{T}(\alpha) \rightarrow \beta$ whose meaning is that "normally, if α then β", or "in the typical situations when α holds, β also holds". They correspond to conditional implications $\alpha \mathrel{|\!\sim} \beta$ of KLM preferential logics [3]. As in PTL [19], the typicality operator cannot be nested. When α and β do not contain occurrences of the typicality operator, an implication $\alpha \rightarrow \beta$ is called *strict*. In the language $\mathcal{L}^{\mathbf{T}}$, we allow *general implications* $\alpha \rightarrow \beta$, where α and β may contain occurrences of the typicality operator.

The interpretation of a typicality formula $\mathbf{T}(\alpha)$ is defined with respect to a preferential interpretation $I = (\mathcal{D}, \Sigma)$ with Σ well-founded.

Definition 3. *Given a preferential interpretation $I = (\mathcal{D}, \Sigma)$, and a labelling $\sigma \in \Sigma$, the valuation of a propositional formula $\mathbf{T}(\alpha)$ in σ is defined as follows:*

$$\sigma(\mathbf{T}(\alpha)) = \begin{cases} \sigma(\alpha) & \textit{if } \sigma \in min_{<_\alpha}(\Sigma) \\ 0_{\mathcal{D}} & \textit{otherwise} \end{cases} \tag{1}$$

where $min_{<_\alpha}(\Sigma) = \{\sigma : \sigma \in \Sigma$ *and* $\nexists \sigma' \in \Sigma$ *s.t.* $\sigma' <_\alpha \sigma\}$.

When $\sigma(\mathbf{T}(A)) > 0_\mathcal{D}$, σ is a labelling assigning a maximal degree of acceptability to argument A in I, i.e., it maximizes the acceptability of argument A, among all the labellings in I.

Given a preferential interpretation $I = (\mathcal{D}, \Sigma)$, we can now define the satisfiability in I of a *graded implication*, having form $\alpha \rightarrow \beta \geq l$ or $\alpha \rightarrow \beta \leq u$, with l and u in $\mathcal{D}$ and α and β boolean combination of arguments. We first define the truth degree of an implication $\alpha \rightarrow \beta$ wrt a preferential interpretation I as follows:

Definition 4. *Given a preferential interpretation* $I = (\mathcal{D}, \Sigma)$ *of an argumentation graph G, the* truth degree of an implication $\alpha \rightarrow \beta$ wrt. I *is defined as:*

$$(\alpha \rightarrow \beta)^I = inf_{\sigma \in \Sigma}(\sigma(\alpha) \rhd \sigma(\beta)).$$

As a special case, for conditional implications, we have that: $(\mathbf{T}(\alpha) \rightarrow \beta)^I = inf_{\sigma \in \Sigma}(\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\beta))$.

We can now define the satisfiability of a *graded implication* in an interpretation $I = (\mathcal{D}, \Sigma)$.

Definition 5. *Given a preferential interpretation* $I = (\mathcal{D}, \Sigma)$ *of an argumentation graph G,* I satisfies a graded implication $\alpha \rightarrow \beta \geq l$ *(written* $I \models \alpha \rightarrow \beta \geq l$*) iff* $(\alpha \rightarrow \beta)^I \geq l$*;* I *satisfies a graded implication* $\alpha \rightarrow \beta \leq u$ *(written* $I \models \alpha \rightarrow \beta \leq u$*) iff* $(\alpha \rightarrow \beta)^I \leq u$.

Example 2 ([18])**.** *As mentioned before, for the weighted argumentation graph in Figure 1, there are 36 labellings in case* $n = 5$*. The following graded conditionals are among the ones satisfied in the interpretation:*

$\mathbf{T}(A_1 \wedge A_2 \wedge \neg A_3) \rightarrow A_6 \geq 1$
(with 4 preferred labellings);
$\mathbf{T}(A_1 \wedge A_2) \rightarrow A_6 \geq 4/5$ *(12 preferred labellings);*
$\mathbf{T}(A_6) \rightarrow A_1 \wedge A_2 \geq 4/5$ *(1 preferred labelling).*

On the other hand, for instance, the strict implication $A_6 \rightarrow A_1 \wedge A_2 \geq 1/5$ *does not hold.*

For instance, this means that the strict implication $A_6 \rightarrow A_1 \wedge A_2$ has a very low degree but, in the situations (labellings) which maximize the acceptability of argument A_6, implication $A_6 \rightarrow A_1 \wedge A_2$ holds with a degree not lower that $4/5$, i.e., roughly speaking, in the labellings maximizing the acceptability of argument A_6, arguments A_1 and A_2 are likely to hold.

Notice that the valuation of a graded implication (e.g., $\alpha \rightarrow \beta \geq l$) in a preferential interpretation I is two-valued, that is, either the graded implication is satisfied in I (i.e., $I \models \alpha \rightarrow \beta \geq l$) or it is not (i.e., $I \not\models \alpha \rightarrow \beta \geq l$). Hence, it is natural to consider *boolean combinations of graded implications*, such as $(\mathbf{T}(A_1) \rightarrow A_2 \wedge A_3 \leq 0.7) \wedge (\mathbf{T}(A_3) \rightarrow A_4) \geq 0.6) \rightarrow (\mathbf{T}(A_1) \rightarrow A_4) \geq 0.6)$, and define their satisfiability in an interpretation I in the obvious way, based on the semantics of classical propositional logic.

The preferential interpretation I^S_G can be used to validate (under the semantics S) properties of interest of an argumentation graph G, expressed by graded implications (including strict or defeasible implications or their boolean combination) based on the semantics S. When the preferential interpretation I^S_G is finite (contains a finite set of labellings), the satisfiability of graded implications (or their boolean combinations) can be verified by *model checking* over the preferential interpretation I^S_G (see section 4). In case there are infinitely many labellings of the graph in the semantics S, which may give rise to non well-founded preference relations, approximations of the semantics S over a finite domain can be considered for proving properties of the argumentation graph. In [18] we have developed an ASP approach for defeasible reasoning over an argumentation graph under the φ-coherent semantics in the finitely-valued case.

3. KLM properties of conditionals

In this section we reformulate the KLM properties of a *preferential consequence relation* in the many-valued setting and prove that, for the choice of combination functions as in Gödel logic, they are satisfied by the set of graded conditionals of the form $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, which hold in a given interpretation $I^S = (\mathcal{S}, \Sigma^S)$.

The KLM postulates of a preferential consequence relations [21, 3, 33] can be reformulated by replacing a conditional $\alpha \mathrel{|\!\sim} \beta$ with the conditional implication $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, as follows:

$(\mathbf{Reflexivity})$ $\mathbf{T}(\alpha) \rightarrow \alpha \geq 1$
$(\mathbf{LeftLogicalEquivalence})$ If $\models \alpha \leftrightarrow \beta$ and $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$, then $\mathbf{T}(\beta) \rightarrow \gamma \geq 1$
$(\mathbf{RightWeakening})$ If $\models \gamma \rightarrow \beta$ and $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$, then $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$
$(\mathbf{And})$ If $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ and $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, then $\mathbf{T}(\alpha) \rightarrow \gamma \wedge \beta \geq 1$
$(\mathbf{Or})$ If $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ and $\mathbf{T}(\beta) \rightarrow \gamma \geq 1$, then $\mathbf{T}(\alpha \vee \beta) \rightarrow \gamma \geq 1$
$(\mathbf{CautiousMonotonicity})$ If $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ and $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, then $\mathbf{T}(\alpha \wedge \beta) \rightarrow \gamma \geq 1$.

Here, we also reinterpret $\models \alpha \rightarrow \beta$ as the requirement that $\alpha \rightarrow \beta \geq 1$ is satisfied in all interpretations $I = (\mathcal{S}, \Sigma)$, that is, $\sigma(\alpha) \rhd \sigma(\beta) \geq 1$ holds for any labelling $\sigma \in \Sigma$, in any interpretation $I = (\mathcal{S}, \Sigma)$. $\models A \leftrightarrow B$ is interpreted as $\models \alpha \rightarrow \beta$ and $\models \beta \rightarrow \alpha$

Concerning the meaning of the postulates in this context, for instance, the meaning of (And) is that, if $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ and $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$ are both satisfied in I^S, then $\mathbf{T}(\alpha) \rightarrow \gamma \wedge \beta \geq 1$ is also satisfied in I^S. The meaning of $(RightWeakening)$ is that, if it holds that $\models \gamma \rightarrow \beta$ (i.e., $\sigma(\gamma) \rhd \sigma(\beta) \geq 1$ for any labelling σ in any interpretation I), and $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ is satisfied in I^S then $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$ is also satisfied in I^S.

Given the interpretation $I^S = (\mathcal{S}, \Sigma^S)$, associated with an argumentation semantics S of a graph G, we can prove the following result.

Proposition 1. *Under the choice of combination functions as in Gödel logic, interpretation I^S satisfies the KLM postulates of a preferential consequence relation given above.*

Proof (Sketch). Let $I^S = (\mathcal{S}, \Sigma^S)$ be the interpretation associated with an argumentation semantics S of a graph G where the t-norm, s-norm, implication function and negation functions are as in Gödel logic (i.e., $a \otimes b = min\{a, b\}$, $a \oplus b = max\{a, b\}$, $a \rhd b = 1$ *if* $a \leq b$ *and* b *otherwise*; and $\ominus a = 1$ *if* $a = 0$ *and* 1 *otherwise*). We proceed by cases.

$(\mathbf{Reflexivity})$ To prove that $\mathbf{T}(\alpha) \rightarrow \alpha \geq 1$ is satisfied in I^S, we have to prove that $inf_{\sigma \in \Sigma^S}\ \sigma(\mathbf{T}(\alpha)) \rhd \sigma(\alpha) \geq 1$. Let us prove that for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\alpha) \geq 1$.

We consider two cases: $\sigma(\mathbf{T}(\alpha)) = 0$, and $\sigma(\mathbf{T}(\alpha)) > 0$.

If $\sigma(\mathbf{T}(\alpha)) = 0$, $\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\alpha) = 0 \rhd \sigma(\alpha) = 1$, and the thesis holds trivially.

If $\sigma(\mathbf{T}(\alpha)) > 0$, by definition $\sigma(\mathbf{T}(\alpha)) = \sigma(\alpha)$. Hence, $\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\alpha) = 1$, and the thesis holds.

$(\mathbf{RightWeakening})$ Assume $\models \gamma \rightarrow \beta$ holds, i.e., $\sigma(\gamma) \rhd \sigma(\beta) \geq 1$ holds for all labellings σ in any preferential interpretation $I = (\mathcal{S}, \Sigma)$. Hence, $inf_{\sigma \in \Sigma}\ \sigma(\gamma) \rhd \sigma(\beta) \geq 1$ and, for all $\sigma \in \Sigma$, $\sigma(\gamma) \rhd \sigma(\beta) \geq 1$. This implies that, for all $I = (\mathcal{S}, \Sigma)$, and for all $\sigma \in \Sigma$, $\sigma(\gamma) \leq \sigma(\beta)$ (in particular, this must hold for all σ in Σ^S).

Let us assume that $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ is satisfied in I^S, i.e., $inf_{\sigma \in \Sigma^S}\ \sigma(\mathbf{T}(\alpha)) \rhd \sigma(\gamma) \geq 1$ holds. Hence, for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\gamma) \geq 1$ holds. Thus, for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha)) \leq \sigma(\gamma)$ and, then, $\sigma(\mathbf{T}(\alpha)) \leq \sigma(\gamma) \leq \sigma(\beta)$. From this it follows that $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$ is satisfied in I^S.

$(\mathbf{Or})$ Let us assume that $\mathbf{T}(\alpha) \rightarrow \gamma \geq 1$ and $\mathbf{T}(\beta) \rightarrow \gamma \geq 1$ are satisfied in I^S. Then, $inf_{\sigma \in \Sigma^S}\ \sigma(\mathbf{T}(\alpha)) \rhd \sigma(\gamma) \geq 1$ and $inf_{\sigma \in \Sigma^S}\ \sigma(\mathbf{T}(\beta)) \rhd \sigma(\gamma) \geq 1$ hold. Hence, for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha)) \rhd \sigma(\gamma) \geq 1$ and $\sigma(\mathbf{T}(\beta)) \rhd \sigma(\gamma) \geq 1$ also hold.

As we have seen above, this implies that: for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha)) \leq \sigma(\gamma)$ and $\sigma(\mathbf{T}(\beta)) \leq \sigma(\gamma)$.

To prove that $\mathbf{T}(\alpha \vee \beta) \rightarrow \gamma \geq 1$ is satisfied in I^S, we prove that for all $\sigma \in \Sigma^S$, $\sigma(\mathbf{T}(\alpha \vee \beta)) \leq \sigma(\gamma)$.

If $\sigma(\mathbf{T}(\alpha \vee \beta)) = 0$, the thesis follows trivially.

If $\sigma(\mathbf{T}(\alpha \vee \beta)) > 0$, σ maximizes the acceptability degree for $\alpha \vee \beta$, and there is no $\sigma' \in \Sigma^S$ such that $\sigma'(\mathbf{T}(\alpha \vee \beta)) > \sigma(\mathbf{T}(\alpha \vee \beta))$.

Given that $\sigma(\mathbf{T}(\alpha \vee \beta)) = \sigma(\alpha \vee \beta) = max\{\sigma(\alpha), \sigma(\beta)\}$, it follows that there is no $\sigma' \in \Sigma^S$ such that $max\{\sigma'(\alpha), \sigma'(\beta)\} > max\{\sigma(\alpha), \sigma(\beta)\}$.

Let us assume, without loss of generality, that $max\{\sigma(\alpha), \sigma(\beta)\} = \sigma(\beta)$. Then, there cannot be a $\sigma' \in \Sigma^S$ such that $\sigma'(\beta) > \sigma(\beta)$, that is, σ maximizes the acceptability degree for β.

Furthermore, $\sigma(\mathbf{T}(\beta)) = \sigma(\beta) = \sigma(\mathbf{T}(\alpha \vee \beta))$

From the hypothesis, we know that $\mathbf{T}(\beta) \rightarrow \gamma \geq 1$ is satisfied in I^S and, hence, $\sigma(\mathbf{T}(\beta)) \leq \sigma(\gamma)$ holds. It follows that $\sigma(\mathbf{T}(\alpha \vee \beta)) = \sigma(\mathbf{T}(\beta)) \leq \sigma(\gamma)$, and then $\sigma(\mathbf{T}(\alpha \vee \beta)) \rhd \sigma(\gamma) \geq 1$.

The case where $max\{A^I(x), B^I(x)\} = A(x)$ is similar, and this concludes the case for $\sigma(\mathbf{T}(\alpha \vee \beta)) > 0$.

For the other postulates the proof is similar. $\square$

The KLM properties considered above do not depend on the choice of the negation function. The same properties also hold for Zadeh's logic. However, some of the properties above might not hold depending on other choices of combination functions. Note that whether the KLM properties are intended or not, may depend on the kind of conditionals and on the kind of reasoning one aims at, and it is still a matter of debate [34, 35, 36, 37].

4. Model checking over finite interpretations

In this section we show that, for a *finite interpretation* $I^S = (\mathcal{S}, \Sigma^S)$ associated with an argumentation semantics S of an argumentation graph G, the satisfiability of a conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ in I^S can be decided in polynomial time in the size of Σ^S times the size of the formula $\mathbf{T}(\alpha) \rightarrow \beta$.

To verify the satisfiability of a graded conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ in a preferential interpretation $I^S = (\mathcal{S}, \Sigma^S)$ of the graph G, one has to check that for all labellings $\sigma \in \Sigma^S$, it holds that $\mathbf{T}(\sigma(\alpha)) \rhd \sigma(\beta) \geq k$. In particular, one has to identify all the labellings $\sigma \in \Sigma^S$ which maximize the acceptability degree of the boolean combination of arguments α (i.e., those such that $\sigma(\mathbf{T}(\alpha)) = v > 0$) as, for all other labellings, $\mathbf{T}(\sigma(\alpha)) = 0$ and $0 \rhd \sigma(\beta) \geq k$ holds trivially.

Let $|\Sigma^S|$ be the size of Σ^S. Identifying the labellings which maximize the acceptability degree of α, requires to evaluate the acceptability degree $\sigma'(\alpha)$ of α, for any $\sigma' \in \Sigma^S$. For a given labelling σ' this evaluation is polynomial in $|\alpha|$, the size of α (the number of subformulas of α is polynomial in $|\alpha|$). Then, determinimg the acceptability degree of α for all labellings in Σ^S, requires a polynomial number of steps in $|\Sigma^S| \times |\alpha|$. In particular, a single scan of the list of the labellings in Σ^S, also allows to identify the labellings maximizing the acceptability degree of α (call them $\sigma_1, \ldots, \sigma_h$), the value $v = \sigma_1(\alpha) = \ldots =$

$\sigma_h(\alpha)$, and the acceptability degree of β in each labelling. Overall this requires a polynomial number of steps in $|\Sigma^S| \times (|\alpha| + |\beta|)$.

Considering that $v = \mathbf{T}(\sigma_1(\alpha)) = \ldots = \mathbf{T}(\sigma_h(\alpha))$, the verification that $v \rhd \sigma_i(\beta) \geq k$ holds, for all $i = 1, \ldots, h$, may require in the worst case a polynomial number of steps in $|\Sigma^S| \times |\beta|$. Overall, the following proposition holds.

Proposition 2. *Given a* finite *interpretation* $I^S = (\mathcal{S}, \Sigma^S)$*, associated with an argumentation semantics* S *of a graph* G*. The satisfiability of a graded conditional* $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ *in* I^S *can be decided in* $O(|\Sigma^S| \times (|\alpha| + |\beta|))$.

Of course, whether the interpretation I^S is finite or not depends on the argumentation semantics S under consideration. For the finitely-valued φ-coherent semantics, an ASP based approach for the verification of graded conditionals has been considered in [18], through a mapping of an argumentation graph to a weighted knowledge base, for which ASP encodings have been developed [38].

5. Towards a probabilistic semantics for gradual argumentation

When the domain of argument valuation is the interval $[0, 1]$, the definition of a preferential interpretation $I^S = (\mathcal{D}, \Sigma^S)$ associated with the gradual semantics S of an argumentation graph G, which has been developed in Section 2, also suggests a probabilistic argumentation semantics, inspired to Zadeh's *probability of fuzzy events* [24]. The approach has been previously considered in [39] for providing a probabilistic interpretation of Self-Organising Maps [40] after training, by exploiting a recent characterization of the continuous t-norms compatible with Zadeh's probability of fuzzy events (P_Z-compatible t-norms) by Montes et al. [41]. In this section we explore this approach in the context of gradual argumentation, to see that it leads to a generalization of the probabilistic semantics presented in [25], and we discuss some advantages and drawbacks of the approach.

Let Σ be the set of labellings of G in a gradual argumentation semantics S with domain of argument valuation in $[0, 1]$, and I_G^S the associated preferential interpretation. The probabilistic semantics we propose is inspired to Zadeh's probability of fuzzy events [24], as one can regard an argument $A \in \mathcal{A}$ as a *fuzzy event*, with *membership function* $\mu_A : \Sigma^S \rightarrow [0, 1]$, where $\mu_A(\sigma) = \sigma(A)$. Similarly, any boolean combination of arguments α can as well be regarded as a fuzzy event, with membership function $\mu_\alpha(\sigma) = \sigma(\alpha)$, where the extension of labellings to boolean combinations of arguments and to typicality formulas has been defined in Section 2.

We restrict to a P_Z-compatible t-norm $\otimes$ [41], with associated t-conorm $\oplus$ and the negation function $\ominus x = 1 - x$. For instance, one can take the minimum t-norm, product t-norm, or Lukasiewicz t-norm. Given $I_G^S = \langle \mathcal{D}, \Sigma^S \rangle$, we assume a discrete probability distribution $p : \Sigma^S \rightarrow [0, 1]$ over Σ^S, and define the *probability of a boolean combination of arguments* α as follows:

$$P(\alpha) = \sum_{\sigma \in \Sigma^S} \sigma(\alpha)\, p(\sigma) \tag{2}$$

For a single argument $A \in \mathcal{A}$, when labellings are two-valued (that is, $\sigma(A)$ is 0 or 1), the definition above becomes the following: $P(A) = \sum_{\sigma \in \Sigma^S \wedge \sigma(A)=1} p(\sigma)$, which relates to the probability of an argument in the probabilistic semantics by Thimm in [25]. Indeed, in [25] the probability of an argument A in Arg is "the degree of belief that A is in an extension", defined as the sum of the probabilities of all possible extensions e that contain argument A, i.e., $P(A) = \sum_{A \in e \subseteq Arg} p(e)$, where an extension $e \in 2^{Arg}$ is a set of arguments in Arg, and $p(e)$ is the probability that e is an extension.

In Thimm's semantics [25] a notion of p-justifiable probability function is introduced to restrict to the "probability functions that agree with our intuition of argumentation", so that relationships with classical argumentation semantics can be established. Here, instead, for a given gradual semantics S with labellings Σ^S, we only consider probability functions on Σ^S (rather than on the set of all possible labellings over the domain). This forces the adherence to the semantics S.

Following Smets [42], we let the *conditional probability of* α *given* β, where α and β are boolean combinations of arguments, to be defined as $P(\alpha|\beta) = P(\alpha \wedge \beta)/P(\beta)$ (provided $P(\beta) > 0$). As observed by Dubois and Prade [43], this generalizes both conditional probability and the fuzzy inclusion index advocated by Kosko [44].

Let us extend the language $L^{\mathbf{T}}$ by introducing a new proposition $\{\sigma\}$, for each $\sigma \in \Sigma$.[2] We extend the valuations σ to such propositions by letting: $\sigma(\{\sigma\}) = 1$ and $\sigma'(\{\sigma\}) = 0$, for any $\sigma' \in \Sigma$ such that $\sigma' \neq \sigma$. It can be proven (see [39]) that

$$P(A|\{\sigma\}) = \sigma(A).$$

The result holds when the t-norm is chosen as in Gödel, Łukasiewicz or Product logic. In such cases, $\sigma(A)$ can be interpreted as the conditional probability that argument A holds, given labelling σ, which can be regarded as a *subjective probability* (i.e., the degree of belief we put into A when we are in a state represented by labelling σ).

Under the assumption that the probability distribution p is *uniform* over the set Σ of labellings, it holds that

[2] A proposition $\{\sigma\}$ corresponds to a *nominal* in description logics [45].

$P(\alpha|\beta) = M(\alpha \wedge \beta)/M(\beta)$ (provided $M(\beta) > 0$), where $M(\alpha) = \sum_{x \in \Sigma} \sigma(A)$ is the *size* of the fuzzy event α. For a finite set of labellings $\Sigma^S = \{\sigma_1, \ldots, \sigma_m\}$ wrt. a given semantics S, assuming a uniform probability distribution, we have that $P(\alpha) = M(\alpha)/m = (\sigma_1(\alpha) + \ldots + \sigma_m(\alpha))/m$. We refer to [18] for an example referring to the φ-coherent semantics.

The notion of probability P defined by equation (2) satisfies Kolmogorov's axioms for any P_Z-compatible t-norm, with associated t-conorm, and the negation function $\ominus x = 1 - x$ [41]. However, there are properties of classical probability which do not hold (depending on the choice of t-norm), as a consequence of the fact that not all classical logic equivalences hold in a fuzzy logic.

For instance, the truth degree of $A \wedge \neg A$ in a labelling σ may be different from 0 depending on the t-norm (e.g., with Gödel and Product t-norms). Hence, it may be the case that $P(A \wedge \neg A)$ is different from 0. Similarly, it may be the case that $P(A \vee \neg A)$ is different from 1 (e.g., with Göedel t-norm) and that $P(A|A) = P(A \wedge A)/P(A)$ is different from 1 (e.g., with Product t-norm). While $P(A) + P(\neg A) = 1$ holds (due to the choice of negation function), $P(A|B) + P(\neg A|B)$ may be different from 1.

While this approach can be regarded as a simple generalization of probabilistic semantics to the gradual case, on the negative side, some properties of classical probability are lost. Hence, we can consider this approach only as a first step towards a probabilistic semantics for gradual argumentation.

6. Conclusions

In this paper we have developed a general approach to define a many-valued preferential interpretation of an argumentation graph, based on a gradual argumentation semantics (i.e., a set of many-valued labellings). The approach allows for graded (strict or conditional) implications involving arguments and boolean combination of arguments (with typicality) to be evaluated in the preferential interpretation I^S of the argumentation graph, which can be defined based on a given gradual argumentation semantics S. We have proven that graded conditionals of the form $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, which are satisfied in I^S, satisfy the postulates of a preferential consequence relation [21], under some choice of combination functions. When the preferential interpretation I^S is finite, the validation of graded conditionals can be done by model-checking over interpretation I^S. The satisfiability of a conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ in I^S can be decided in polynomial time in the size of Σ^S times the size of the formula $\mathbf{T}(\alpha) \rightarrow \beta$. For the gradual semantics with domain of argument valuation in the unit real interval $[0, 1]$, the paper also proposes a probabilistic argumentation semantics, which builds on a gradual semantics S and on the preferential interpretation I^S of G, and is inspired by Zadeh's probability of fuzzy events [24].

Concerning the relationships between argumentation semantics and conditional reasoning, Weydert [7] has proposed one of the first approaches for combining abstract argumentation with a conditional semantics. He has studied "how to interpret abstract argumentation frameworks by instantiating the arguments and characterizing the attacks with suitable sets of conditionals describing constraints over ranking models". In doing this, he has exploited the JZ-evaluation semantics, which is based on system JZ [46]. Our approach does not commit to any specific gradual argumentation semantics, and aims at providing a preferential and conditional interpretation for a large class of gradual argumentation semantics.

For Abstract Dialectical Frameworks (ADFs) [8], the correspondence between ADFs and Nonmonotonic Conditional Logics has been studied in [9] with respect to the two-valued models, the stable, the preferred semantics and the grounded semantics of ADFs.

In [10] Ordinal Conditional Functions (OCFs) are interpreted and formalized for Abstract Argumentation, by developing a framework that allows to rank sets of arguments with respect to their plausibility. An attack from argument a to argument b is interpreted as the conditional relationship, "if a is acceptable then b should not be acceptable". Based on this interpretation, an OCF inspired by System Z ranking function is defined. In this paper we focus on the gradual case, based on a many-valued logic.

In [29, 30] an approach is presented which regards a weighted argumentation graph as a weighted conditional knowledge base in a fuzzy defeasible Description Logic. In this approach, a pair of arguments $(B, A) \in \mathcal{R}$ with weight w_{AB} (representing an attack or a support), corresponds to a conditional implication $\mathbf{T}(A) \sqsubseteq B$ with weight w_{AB}. Based on this correspondence, some semantics for weighted knowledge bases with typicality [47] have inspired some argumentation semantics [29], and vice-versa. In particular, in [18] we have developed an ASP approach for defeasible reasoning over an argumentation graph under the φ-coherent semantics in the finitely-valued case. In this paper, we have generalized the approach beyond the φ-coherent semantics, to deal with a large class of gradual semantics.

In Section 5, we have proposed a probabilistic semantics for gradual argumentation, which builds on the many-valued interpretation of the argumentation graph, and is inspired to Zadeh's *probability of fuzzy events* [24]. Under this semantics, the truth degree $\sigma(A)$ of an argument A in a labelling σ can be regarded as the conditional probability of A given σ. The proposed approach can be seen as a generalization of the probabilistic semantics by Thimm [25] to the gradual case, but with some differences. On the one hand, our approach does not require to introduce a notion of p-justifiable probability function, as it only considers probability functions on the set of labellings Σ^S of a graph in a given gradual semantics S. On the other hand, as we have seen, some classical equivalences may not hold (depending on the choice of combination functions), and some properties of classical probability may be lost. This requires further investigation. Alternative approaches for combining conditionals and probabilities, such as the one recently proposed by Flaminio et al. [48], might suggest alternative ways of defining a probabilistic semantics for gradual argumentation.

In this paper, we have adopted an epistemic approach to probabilistic argumentation, and we refer to [49] for a general survey on probabilistic argumentation. As a generalization of the epistemic approach to probabilistic argumentation, epistemic graphs [31] allow for epistemic constraints, that is, for boolean combinations of inequalities, involving statements about probabilities of formulae built out of arguments. While the conditional many-valued semantics in Section 2 allows for combining graded conditionals, we have not considered graded conditionals in the probabilistic semantics. This might be a possible direction to extend the probabilistic semantics in Section 5.

Acknowledgements: We thank the anonymous referees for their helpful suggestions. This research was partially supported by INDAM-GNCS. Mario Alviano was partially supported by Italian Ministry of Research (MUR) under PNRR project FAIR "Future AI Research", CUP H23C22000860006 and by LAIA lab (part of the SILA labs).

References

[1] Q. Zhong, X. Fan, X. Luo, F. Toni, An explainable multi-attribute decision model based on argumentation, Expert Syst. Appl. 117 (2019) 42–61.

[2] J. Pearl, System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: TARK'90, Pacific Grove, CA, USA, 1990, pp. 121–135.

[3] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial Intelligence 55 (1992) 1–60.

[4] H. Geffner, J. Pearl, Conditional entailment: Bridging two approaches to default reasoning, Artif. Intell. 53 (1992) 209–244.

[5] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: Proc. IJCAI'93, Chambéry,, 1993, pp. 640–647.

[6] P. M. Dung, On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games, Artif. Intell. 77 (1995) 321–357.

[7] E. Weydert, On the plausibility of abstract arguments, in: Proc. Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 12th European Conference, ECSQARU 2013, Utrecht, The Netherlands, July 8-10, 2013, volume 7958 of *LNCS*, Springer, 2013, pp. 522–533.

[8] G. Brewka, H. Strass, S. Ellmauthaler, J. P. Wallner, S. Woltran, Abstract dialectical frameworks revisited, in: IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013, IJCAI/AAAI, 2013, pp. 803–809.

[9] J. Heyninck, G. Kern-Isberner, M. Thimm, On the correspondence between abstract dialectical frameworks and nonmonotonic conditional logics, in: Proceedings of the Thirty-Third International Florida Artificial Intelligence Research Society Conference, May 17-20, 2020, AAAI Press, 2020, pp. 575–580.

[10] K. Skiba, M. Thimm, Ordinal conditional functions for abstract argumentation, in: COMMA 2022, Cardiff, Wales, UK, 14-16 September 2022, volume 353 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2022, pp. 308–319.

[11] C. Cayrol, M. Lagasquie-Schiex, Graduality in argumentation, J. Artif. Intell. Res. 23 (2005) 245–297. doi:10.1613/jair.1411.

[12] J. Janssen, M. D. Cock, D. Vermeir, Fuzzy argumentation frameworks., in: IPMU 2008, 2008, pp. 513–520.

[13] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, M. J. Wooldridge, Weighted argument systems: Basic definitions, algorithms, and complexity results, Artif. Intell. 175 (2011) 457–486.

[14] S. Egilmez, J. G. Martins, J. Leite, Extending social abstract argumentation with votes on attacks, in: TAFA 2013, Beijing, China, Aug. 3-5, LNCS 8306, Springer, 2013, pp. 16–31.

[15] L. Amgoud, J. Ben-Naim, D. Doder, S. Vesic, Acceptability semantics for weighted argumentation frameworks, in: IJCAI 2017, Melbourne, Australia, 2017, pp. 56–62.

[16] P. Baroni, A. Rago, F. Toni, How many properties do we need for gradual argumentation?, in: Proc. AAAI 2018, New Orleans, Louisiana, USA, February 2-7, 2018, pp. 1736–1743.

[17] L. Amgoud, D. Doder, Gradual semantics accounting for varied-strength attacks, in: Proceedings AAMAS '19, Montreal, QC, Canada, May 13-17, 2019, 2019, pp. 1270–1278.

[18] M. Alviano, L. Giordano, D. Theseider Dupré, Preferential reasoning with typicality in ASP over weighted argumentation graphs in a gradual semantics, in: Proc. ICLP Workshops 2023, CEUR Workshop Proceedings 3437, 2023.

[19] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On

rational entailment for propositional typicality logic, Artif. Intell. 277 (2019).

[20] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, ALC+T: a preferential extension of Description Logics, Fundamenta Informaticae 96 (2009) 1–32.

[21] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial Intelligence 44 (1990) 167–207.

[22] T. Lukasiewicz, U. Straccia, Managing uncertainty and vagueness in description logics for the semantic web, J. Web Semant. 6 (2008) 291–308.

[23] L. Giordano, D. Theseider Dupré, An ASP approach for reasoning on neural networks under a finitely many-valued semantics for weighted conditional knowledge bases, Theory Pract. Log. Program. 22 (2022) 589–605.

[24] L. Zadeh, Probability measures of fuzzy events, J.Math.Anal.Appl 23 (1968) 421–427.

[25] M. Thimm, A probabilistic semantics for abstract argumentation, in: 20th European Conf. on Art. Intell., Montpellier, France, August 27-31 , ECAI 2012, volume 242 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2012, pp. 750–755.

[26] P. Baroni, A. Rago, F. Toni, From fine-grained properties to broad principles for gradual argumentation: A principled spectrum, Int. J. Approx. Reason. 105 (2019) 252–286.

[27] N. Potyka, Interpreting neural networks as quantitative argumentation frameworks, in: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, February 2-9, 2021, AAAI Press, 2021, pp. 6463–6470.

[28] L. Amgoud, C. Cayrol, M. Lagasquie-Schiex, On the bipolarity in argumentation frameworks, in: 10th International Workshop on Non-Monotonic Reasoning (NMR 2004), Whistler, Canada, June 6-8, 2004, Proceedings, 2004, pp. 1–9.

[29] L. Giordano, From weighted conditionals of multilayer perceptrons to a gradual argumentation semantics, in: 5th Workshop on Advances in Argumentation in Artif. Intell., 2021, Milan, Italy, Nov. 29, volume 3086 of *CEUR Workshop Proc.*, 2021. Extended version in CoRR abs/2110.03643.

[30] L. Giordano, From weighted conditionals with typicality to a gradual argumentation semantics and back, in: 20th International Workshop on Non-Monotonic Reasoning, NMR 2022, Part of FLoC 2022, Haifa, Israel, August 7-9, 2022, volume 3197 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 127–138.

[31] A. Hunter, S. Polberg, M. Thimm, Epistemic graphs for representing and reasoning with positive and negative influences of arguments, Artif. Intell. 281 (2020) 103236.

[32] S. Gottwald, A Treatise on Many-valued Logics, Research Studies Press, 2001.

[33] J. Pearl, Probabilistic Reasoning in Intelligent Systems Networks of Plausible Inference, Morgan Kaufmann, 1988.

[34] P. A. Bonatti, L. Sauro, On the logical properties of the nonmonotonic description logic DL^{N}, Artif. Intell. 248 (2017) 85–111.

[35] C. D. Koutras, K. Liaskos, C. Moyzes, C. Rantsoudis, Default reasoning via topology and mathematical analysis: A preliminary report, in: Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018, 2018, pp. 267–276.

[36] H. Rott, Difference-making conditionals and the relevant ramsey test, The Review of Symbolic Logic 15 (2019).

[37] G. Casini, T. Meyer, I. Varzinczak, Taking defeasible entailment beyond rational closure, in: Logics in Artificial Intelligence - 16th European Conference, JELIA 2019, Rende, Italy, May 7-11, 2019, Proceedings, 2019, pp. 182–197.

[38] M. Alviano, L. Giordano, D. Theseider Dupré, Complexity and scalability of defeasible reasoning in many-valued weighted knowledge bases, CoRR abs/2303.04534 (2023). URL: https://doi.org/10.48550/arXiv.2303.04534.

[39] L. Giordano, V. Gliozzi, D. Theseider Dupré, A conditional, a fuzzy and a probabilistic interpretation of self-organizing maps, J. Log. Comput. 32 (2022) 178–205.

[40] T. Kohonen, M. Schroeder, T. Huang (Eds.), Self-Organizing Maps, Third Edition, Springer Series in Information Sciences, Springer, 2001.

[41] I. Montes, J. Hernández, D. Martinetti, S. Montes, Characterization of continuous t-norms compatible with zadeh's probability of fuzzy events, Fuzzy Sets Syst. 228 (2013) 29–43.

[42] P. Smets, Probability of a fuzzy event: An axiomatic approach, Fuzzy Sets and Systems 7 (1982) 153–164.

[43] D. Dubois, H. Prade, Fuzzy sets and probability: misunderstandings, bridges and gaps, in: [Proceedings 1993] Second IEEE International Conference on Fuzzy Systems, 1993, pp. 1059–1068 vol.2. doi:10.1109/FUZZY.1993.327367.

[44] B. Kosko, Neural networks and fuzzy systems: a dynamical systems approach to machine intelligence, Prentice Hall, 1992.

[45] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider, The Description Logic Handbook - Theory, Implementation, and Applications, Cambridge, 2007.

[46] E. Weydert, System JLZ - rational default reason-

ing by minimal ranking constructions, Journal of Applied Logic 1 (2003) 273–308.

[47] L. Giordano, D. Theseider Dupré, Weighted defeasible knowledge bases and a multipreference semantics for a deep neural network model, in: Proc. JELIA 2021, May 17-20, volume 12678 of *LNCS*, Springer, 2021, pp. 225–242.

[48] T. Flaminio, L. Godo, H. Hosni, Boolean algebras of conditionals, probability and logic, Artif. Intell. 286 (2020) 103347.

[49] A. Hunter, , S. Polberg, N. Potyka, T. Rienstra, M. Thimm, Probabilistic argumentation: A survey, in: D. Gabbay, M. Giacomin, G. Simari (Eds.), Handbook of Formal Argumentation, Volume 2, College Publications, 2021. URL: https://books.google.it/books?id=JUekzgEACAAJ.