Explaining gradual argumentation semantics in a conditional multi-preferential logic with typicality

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 framework to provide 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.

Keywords

Gradual argumentation, Many-valued semantics, Preferential and Conditional reasoning

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].

The paper proposes a general approach to develop a preferential interpretation of an argumentation graph under a gradual semantics, provided some weak conditions on the domain of argument interpretation are satisfied. The approach allows for *conditional reasoning over the argumentation graph*, by formalizing conditional properties of the graph (with respect to the chosen semantics) in a many-valued logic with typicality: a many-valued propositional logic in which arguments play the role of propositional variables and in which a typicality operator is introduced, inspired by the typicality operator proposed in the Propositional Typicality Logic [18] and in Description Logics (DLs) with typicality [19]. The operator allows for the definition of

7th Workshop on Advances in Argumentation in Artificial Intelligence (AI³), November 06–09, 2023, Rome, Italy

*Corresponding author.
† These authors contributed equally.
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)

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 the chosen (gradual) argumentation semantics. They correspond to conditional implications $\alpha \mathrel{|\!\sim} \beta$ in the KLM approach [20, 3].

More precisely, the paper considers *graded implications* of the form $\alpha \rightarrow \beta \geq l$, where α and β can be *boolean combination of arguments* possibly containing occurrences of the typicality operator. In particular, *graded conditionals* of the form $\mathbf{T}(\alpha) \rightarrow \beta \geq l$ have the meaning that "normally argument α implies argument β with degree at least l". They are inspired by graded inclusion axioms in fuzzy DLs [21] and in weighted defeasible DLs knowledge bases [22].

The satisfiability of such implications in the *multi-preferential* interpretation I_G^S of an argumentation graph G (with respect to some given semantics S), exploits multiple preference relations $<_\alpha$ over labellings, each one associated with a boolean combination of arguments α.

In [23] it has been shown that the satisfiability of a graded conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ in a *finite* preferential interpretation I_G^S can be decided in polynomial time in the product of the size of the interpretation I_G^S and the size of the conditional formula. For well-founded preferences, the KLM postulates of a preferential consequence relation, reformulated for graded conditionals, can be proven to be satisfied by the conditionals which hold in the multi-preferential interpretation I_G^S, for some choice of combination functions. In this paper, we consider an extension of the multi-preferential approach in [23] by lifting the well-foundedness restriction on the preference relations.

2. Gradual argumentation semantics: truth degree set and labellings of a graph

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 [24] 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, 25] (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 the unit interval $[0, 1]$ or, in the finitely-valued case (as in [24]), 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 [26], for a generic domain $\mathcal{D}$. As we want to capture both weighted and non-weighted

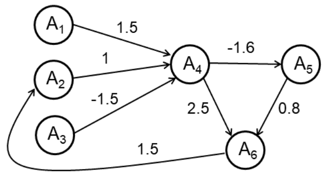

Figure 1: Example weighted argumentation graph G from [24]

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.

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 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 $\pi(B, A) = -1$ of attacks and $\pi(B, A) = +1$ for supports.

Bipolar argumentation has been studied in the literature [27, 16, 25, 26] through different frameworks. We refer to the Quantitative Bipolar Argumentation Framework (QBAF) by Baroni, Rago and Toni [16, 25] for a classification and the properties of gradual semantics, when the argumentation graph is non-weighted, and to Potyka's work [26] for the framework of edge-weighted QBAFs and its properties. The properties of edge-weighted argumentation graphs with weights in $[0, 1]$ have also been studied in Amgoud and Doder's framework [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 $\mathcal{D}$. In some cases, we may omit the base score σ_0, and consider the set of labellings Σ^S of a graph G for all the possible choices of the base score, or for a subset of them. In the following we will assume that, whatever the concrete definition of a semantics S might be, the semantics of G 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.

Example 1 ([24]). *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 labelling for $n = 5$.*

3. A many-valued logic of arguments

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 $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$, let $\mathcal{L}$ be a propositional language whose set of propositional variables $Prop$ is the set of arguments $\mathcal{A}$. We assume that the 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 et al. [28] in their epistemic approach to probabilistic argumentation.

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 $\mathcal{C}_n$, $\otimes$, $\oplus$, $\rhd$ and $\ominus$ can be chosen as a t-norm, an s-norm, an implication function, and a negation function in some system of many-valued logic [29].

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: $\sigma(\alpha \wedge \beta) = \sigma(\alpha) \otimes \sigma(\beta)$, $\sigma(\alpha \vee \beta) = \sigma(\alpha) \oplus \sigma(\beta)$, $\sigma(\alpha \rightarrow \beta) = \sigma(\alpha) \rhd \sigma(\beta)$, and $\sigma(\neg\alpha) = \ominus\sigma(\alpha)$. 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 $\mathcal{L}$ and that $\sigma(\bot) = 0_{\mathcal{D}}$ and $\sigma(\top) = 1_{\mathcal{D}}$, for all labellings σ.

4. A preferential interpretation of an argumentation graph

In this section, given an argumentation graph G and a semantics $(\mathcal{D}, \Sigma^S)$ of G, we aim at defining a preferential interpretation of the graph. We first introduce a preference relation on the set of labellings Σ^S, associated to any boolean combination of arguments α.

Definition 1. *Given a set of labellings Σ^S, for each boolean combination of arguments α , we define a* preference relation $<_\alpha$ *on Σ, as follows: for $\sigma, \sigma' \in \Sigma$, $\sigma <_\alpha \sigma'$ iff $\sigma'(\alpha) < \sigma(\alpha)$.*

Labelling σ *is preferred to* σ' with respect to an argument (or a boolean combination of arguments) α when σ is *more plausible* than σ' for argument α, that is, when the degree of truth of α in σ is greater than the degree of truth of α in σ'. The preference relation $<_\alpha$ is a *strict partial order relation* on Σ.

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.

Let us 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 triple $I^S_G = (\mathcal{D}, \Sigma^S, \{<_\alpha\})$.*

We have explicitly associated the preference relations $<_\alpha$ to the set of labellings Σ^S of the graph, although preference are induced by the labellings in Σ^S. Often, we will simply write I^S or I, rather than I^S_G (and $(\mathcal{D}, \Sigma, \{<_\alpha\})$ rather then $(\mathcal{D}, \Sigma^S, \{<_\alpha\})$).

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*" [18]. The typicality operator allows for 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 [18] and in [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 for *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 $I = (\mathcal{D}, \Sigma, \{<_\alpha\})$.

Definition 3. *Given a preferential interpretation $I = (\mathcal{D}, \Sigma, \{<_\alpha\})$, 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 $\sigma(\mathbf{T}(\alpha)) > 0_{\mathcal{D}}$, σ is a labelling assigning a maximal degree of acceptability to argument α in I, i.e., it maximizes the acceptability of argument α, among all the labellings in I. As we lifted the requirement that preferences $<_\alpha$ are well-founded, the set $min_{<_\alpha}(\Sigma)$ might be empty.

5. Graded implications

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

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

We can now define the satisfiability of a *graded implication* in an interpretation I.

Definition 5. *Given a preferential interpretation $I = (\mathcal{D}, \Sigma, \{<_\alpha\})$ of an argumentation graph G wrt. S, 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 $(\alpha \rightarrow \beta)^I \leq u$.*

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_G^S can be used to validate 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. For instance, the boolean combination of graded conditionals above allows to verify to validate whether the graded conditional $(\mathbf{T}(A_1) \rightarrow A_4) \geq 0.6)$ holds, for all the labellings of the graph G (in the semantics S) satisfying the graded conditionals $(\mathbf{T}(A_1) \rightarrow A_2 \wedge A_3 \leq 0.7)$ and $(\mathbf{T}(A_3) \rightarrow A_4) \geq 0.6)$.

When the preferential interpretation I_G^S is finite (i.e., it contains a finite set of labellings), the satisfiability of graded implications (and their boolean combinations) can be verified by *model checking* over the preferential interpretation I_G^S. In case there are infinitely many labellings of the graph in the semantics S, approximations of the semantics S over a finite domain can be considered for proving properties of the argumentation graph. As a proof of concept, in [24] we have developed an ASP approach for defeasible reasoning over an argumentation graph under the φ-coherent semantics in the finitely-valued case.

6. Related Work

In [7] Weydert 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 [30].

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.

Our approach does not commit to a specific gradual argumentation semantics, and aims at providing a preferential conditional interpretation for a large class of gradual argumentation semantics. In this paper we focus on the gradual case, based on a many-valued logic.

In [31, 32] 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 [33] have inspired some argumentation semantics [31], and vice-versa. In particular, in [24] 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.

7. Conclusions

In this paper, we have developed a general framework to define a many-valued preferential interpretation of an argumentation graph, with respect to a gradual argumentation semantics. The approach allows for graded (strict and conditional) implications involving arguments and boolean combination of arguments (with typicality) to be evaluated in the preferential interpretation I_G^S of the argumentation graph, which can be constructed based on a given gradual argumentation semantics S. When the preferential interpretation I_G^S is finite, the validation of graded conditionals can be done by model-checking over interpretation I_G^S.

In [23] we have shown that graded conditionals $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$, which are satisfied in I_G^S, satisfy the postulates of a preferential consequence relation [20] (suitable reformulated in this setting), for some choice of combination functions. Whether such properties are satisfied by the semantics considered in this paper, which does not require the preference relations to be well-founded, will be a subject of future work.

The definition of a preferential interpretation I_G^S associated with an argumentation graph G and a gradual semantics S also sets the ground for the definition of a *probabilistic interpretation* for gradual semantics with domain of argument valuation in the unit real interval $[0, 1]$. Such interpretation is inspired by Zadeh's probability of fuzzy events [34], and can be regarded as a generalization of the probabilistic semantics by Thimm [35] to the gradual case. We refer to [23] for details.

Acknowledgments

Laura Giordano and Daniele Theseider Dupré were partially supported by the Università del Piemonte Orientale, and by INDAM-GNCS Project 2022 " LESLIE: LogichE non-claSsiche per tooL Intelligenti ed Explainable". 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] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On rational entailment for propositional typicality logic, Artif. Intell. 277 (2019).

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

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

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

[22] 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.

[23] M. Alviano, L. Giordano, D. Theseider Dupré, Typicality, conditionals and a probabilistic semantics for gradual argumentation, in: Proc. 21st International Workshop on Non-Monotonic Reasoning (NMR 2023), Rhodes, Greece, September 2-4, 2023, volume 3464 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 4–13.

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

[25] 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.

[26] 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.

[27] 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.

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

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

[30] E. Weydert, System JLZ - rational default reasoning by minimal ranking constructions, Journal of Applied Logic 1 (2003) 273–308.

[31] 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.

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

[33] 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.

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

[35] 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.