Conditionals and Temporal Conditionals 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 develop a preferential interpretation of gradual argumentation and propose an approach for temporal conditional reasoning about argumentation graphs. The approach exploits a two-valued temporal conditional logic with typicality, combining a preferential logic with Linear Time Temporal Logic (LTL). It introduces a dynamic dimension to conditional reasoning in gradual argumentation, enabling the verification of conditional properties across time, such as trends in the evolution of argument strength.

Keywords
Preferential and Conditional reasoning, Argumentation, Temporal Reasoning

1. Introduction

Conditional reasoning plays a central role in artificial intelligence, especially when distinguishing between strict rules and typical situations. In previous work [1, 2], we have presented a many-valued conditional logic with typicality to provide preferential interpretation of gradual argumentation [3, 4, 5, 6, 7, 8]. This framework enables conditional reasoning over arguments and their boolean combinations, evaluated with respect to some chosen gradual semantics, through the verification of graded (strict or defeasible) implications over an argumentation graph.

In this paper, we aim at exploring a two-valued approach, by exploiting a two-valued conditional logic with typicality to provide a preferential interpretation of gradual argumentation. Then, we extend the approach for dealing with the temporal case. A temporal conditional logic with typicality can be exploited in the verification of the transient behavior of an argumentation graph, as well as in the verification of the transient behavior of a neural network. Indeed, a multilayer (recurrent) neural network can be regarded as a weighted knowledge base [9, 10], and also as a weighted argumentation graph [11, 12, 13].

In a nutshell, our contribution is twofold:

1. *Conceptual simplification:* By moving from a many-valued to a two-valued framework, we retain expressive power while simplifying the formal semantics and reasoning mechanisms.
2. *Temporal extension:* We introduce a dynamic dimension to conditional reasoning in gradual argumentation, enabling the verification of conditional properties across time, such as trends in the evolution of argument strength.

The proposed logic can be used to analyze time-dependent properties of argumentation graphs, including iterative updates of strength functions, and is also applicable to weighted knowledge bases and neural networks interpreted as weighted argumentation frameworks. This opens up new opportunities for explainability, temporal verification, and symbolic-neural integration.

AI[3] 2025: 9th Workshop on Advances in Argumentation in Artificial Intelligence, September 13, 2025, Rende, Italy
*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)

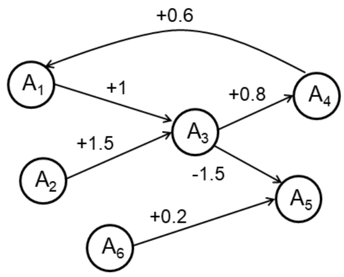

Figure 1: Example weighted argumentation graph G

In more detail, our approach combines preferential approaches to commonsense reasoning [14, 15, 16, 17, 18, 19, 20, 21, 22] with the Linear Time Temporal Logic (LTL) [23]. Preferential extensions of LTL with defeasible temporal operators have been recently studied to enrich temporal formalisms with non-monotonic reasoning features, by considering defeasible versions of the LTL operators [24, 25, 26]. In this regard, we follow a different route, adding the standard LTL operators to a conditional logic with typicality, an approach similar to the one pursued for Description Logics (DLs), through an extension of the temporal description logic $LTL_{\mathcal{ALC}}$ [27] with the typicality operator [28]. As in the Propositional Typicality Logic (PTL) by Booth et al. [29] (and in the DLs with typicality [30]) the conditionals will be formalized based on material implication plus a *typicality operator* $\mathbf{T}$. The typicality operator allows for the definition of *conditional implications* $\mathbf{T}(\alpha) \rightarrow \beta$, meaning that "normally if α holds, β holds". They correspond to conditional implications $\alpha \mid\!\sim \beta$ in KLM logics [17, 19]. In the paper, we will consider a multi-preferential logic, where preferences are associated to aspects (and to arguments).

The structure of the paper is as follows. Section 2 provides the necessary background on gradual argumentation and two-valued multi-preferential conditional logic. In Section 3, we show how such logic can be instantiated for reasoning on argumentation graphs. Section 4 introduces the temporal extension of the logic, and Section 5 describes how it can be instantiated to model temporal aspects of gradual argumentation. Section 6 concludes and discusses future directions.

2. Preliminaries

In this section we provide preliminaries on gradual argumentation semantics and on a two-valued multi-preferential semantics for conditionals.

2.1. Gradual argumentation semantics

In this following, we shortly recap gradual argumentation semantics, following Baroni, Rago and Toni [6, 7], and consider a specific semantic from [31, 32, 13].

As in the Quantitative Bipolar Argumentation Framework (QBAF) by Baroni et al. [6, 7], we let the *domain of argument interpretation* be a set $\mathcal{D}$, equipped with a *preorder relation* $\leq$. In the literature, this assumption is considered general enough to include the domain of argument valuations in most gradual argumentation semantics [3, 33, 4, 34, 5, 6, 8, 11]. We do not assume that $\mathcal{D}$ contains a *minimum element* and a *maximum element*. If they exist, we 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. For the definition of an argumentation graph, we consider the definition of *edge-weighted QBAF* by Potyka [11], for a generic domain $\mathcal{D}$.

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*, the base *score function* $\sigma_0 : \mathcal{A} \rightarrow \mathcal{D}$ assigns a *base score* to arguments, and $\pi : \mathcal{R} \rightarrow \mathbb{R}$ is a *weight function* assigning a positive or negative weight to edges. An example of argumentation graph is in Figure 1, where a base score is not given.

A ir $(B, A) \in \mathcal{R}$ 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.

The properties of edge-weighted argumentation graphs with weights in the interval $[0,1]$ have been studied in [11] as well as in the gradual semantics framework by Amgoud and Doder [8].

Whatever semantics S is considered for an argumentation graph G, we will assume that the semantics S identifies *a set* Σ_G^S *of many-valued labellings* (also called *strength functions*, or *weightings*) of the graph G over a domain of argument valuation $\mathcal{D}$. A *many-valued labelling* σ *of* G *over* $\mathcal{D}$ is a total function $\sigma : \mathcal{A} \rightarrow \mathcal{D}$, which assigns to each argument an *acceptability degree* (or a *strength*) in the domain of argument valuation $\mathcal{D}$.

When, in the following, we want to consider a set of possible values for the initial score of arguments, we will represent the argumentation graph G as a triple $G = (\mathcal{A}, \mathcal{R}, \Sigma_0, \pi)$, where $\Sigma_0 = \{\sigma_0^1, \sigma_0^2, \ldots, \sigma_0^k\}$ is a *finite set of possible initial scores.*

2.2. The φ-coherent semantics

Let us recall the definition of the φ-coherent semantics of an argumentation graph G [13]. We let $\mathcal{D}$ be the interval $[0,1]$ or the finite set $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\}$, for an integer $n \geq 1$. Given a weighted argumentation graph $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$, we let $R^-(A) = \{B \mid (B, A) \in \mathcal{R}\}$. When $R^-(A) = \emptyset$, argument A has neither supports nor attacks.

For a weighted graph $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$ and a many-valued labelling σ, the *weight* W_σ^G *on* $\mathcal{A}$ is defined as a partial function $W_\sigma^G : \mathcal{A} \rightarrow \mathbb{R}$, assigning a positive or negative support (relative to labelling σ) to all arguments $A_i \in \mathcal{A}$ such that $R^-(A_i) \neq \emptyset$, as follows:

$$W_\sigma^G(A_i) = \sum_{A_j \in R^-(A_i)} \pi(A_j, A_i)\ \sigma(A_j) \tag{1}$$

$W_\sigma^G(A_i)$ is left *undefined* when $R^-(A_i) = \emptyset$.

Definition 1. *Given a weighted argumentation graph* $G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$ *and a non-decreasing function* $\varphi : \mathbb{R} \rightarrow \mathcal{D}$*, a* φ-coherent many-valued labelling σ of G *is defined as follows:*

$$\sigma(A_i) = \begin{cases} \varphi(W_\sigma^G(A_i)) & \textit{for all } A_i \in \mathcal{A} \textit{ s.t. } R^-(A_i) \neq \emptyset \\ \sigma_0(A_i) & \textit{otherwise} \end{cases} \tag{2}$$

The semantics is a perceptron-like semantics which has been inspired by some preferential semantics for weighted KBs developed in [35, 10], and has some relations with the semantics proposed by Potyka [11] for interpreting neural networks (see [13] for comparisons). In this perceptron like view, the argumentation graph plays the role of a (possibly recurrent) multilayer network, where arguments A_i correspond to units, edges (with their weights) correspond to synaptic connections between units, $W_\sigma^G(A_i)$ corresponds to the induced local field of unit A_i, and φ corresponds to the activation function. The acceptability degree of an argument A_i in a labelling σ corresponds to the activity of unit A_i in a stationary state of the network. We refer to [9, 13] for further details, including properties of the semantics and comparisons with other argumentation semantics.

We denote by Σ_G the set of all the φ-coherent many-valued labelling σ of $G = \langle \mathcal{A}, \mathcal{R}, \pi \rangle$, for all the possible choices of the initial score; by $\Sigma_G^{\sigma_0}$ the set of the φ-coherent many-valued labelling σ for the initial score σ_0, and by $\Sigma_G^{\Sigma_0}$ the set of the φ-coherent many-valued labelling σ for all the values of initial score in Σ_0.

Example 1. *In the* φ*-coherent semantics for the weighted argumentation graph* G *in Figure 1, in the finitely-valued case with* $\mathcal{D} = \mathcal{C}_n$*, for* $n = 5$*, with* φ *being the approximation in* $\mathcal{C}_n$ *of the logistic function,* Σ_G *contains 36 many-valued* φ*-coherent labellings, while, for* $n = 9$*,* Σ_G *contains 100* φ*-coherent labellings. In this case, there is a labelling for each combination of values of the base score for* A_2 *and* A_6*. For instance,* $\sigma = (3/5, 0, 3/5, 3/5, 1/5, 0)$ *(meaning that* $\sigma(A_1) = 3/5$*,* $\sigma(A_2) = 0$*, and so on) is a many-valued* φ*-coherent labelling of* G *for* $n = 5$*.*

Above, the notion of φ-coherent many-valued labelling of G is defined through a set of equations, as in Gabbay's *equational approach* to argumentation networks [36]. A definition of the φ-coherent semantics can also be given, in the style of the gradual semantics in the framework by Amgoud and Doder [8]. We refer to [13] for details and for a discussion of the properties of the semantics. In particular, one cannot assume that, for an initial score σ_0, there is a unique strength function σ (a unique φ-coherent labelling).

2.3. A two-valued multi-preferential semantics for conditionals

In the following we recall the multi-preferential semantics from [37], and slightly extend it.

As mentioned in the introduction, the conditional logic extends a propositional language L with a typicality operator $\mathbf{T}$, following the approach used in the description logic $\mathcal{ALC} + \mathbf{T}$ [38] as well as in the Propositional Typicality Logic (PTL) [29]. Intuitively, "a sentence of the form $\mathbf{T}(\alpha)$ is understood to refer to the *typical situations in which* α *holds*" [29]. As in PTL [29], the typicality operator cannot be nested. When an implication has the form $\mathbf{T}(\alpha) \rightarrow \beta$, it is called a *defeasible implication*, whose meaning is that "normally, if α then β". An implication $\alpha \rightarrow \beta$ is called *strict*, if it does not contain occurrences of the typicality operator.

The KLM preferential semantics [17, 19, 16] exploits a single preference relation between worlds: *a set of worlds* $\mathcal{W}$, with their valuation and a preference relation $<$ among worlds (where $w < w'$ means that world w is more normal than world w'). A conditional $A \mid\!\sim B$ is satisfied in a KLM preferential interpretation, if B holds in all the most normal worlds satisfying A, i.e., in all $<$-minimal worlds satisfying A. Here, instead, we consider a *multi-preferential semantics*, where preference relations are associated with distinguished propositional formulas $A_1, \ldots, A_m$ (called *distinguished propositions* in the following). In the semantics, a preference relation will be associated with each distinguished proposition A_i, where $w <_{A_i} w'$ means that world w is less atypical than world w' concerning aspect/property A_i (e.g., $w <_{student} w'$ means that w describes a less atypical situation for a student than w').

In the following we will consider *finite* KBs over a set $Prop$ of propositional variables, and a finite set of *distinguished propositions* $A_1, \ldots, A_m$ (propositional formulas over $Prop$). Preferential interpretations are equipped with a set of worlds $\mathcal{W}$ and a finite set of preference relations $<_{A_1}, \ldots, <_{A_n}$, where, for each distinguished proposition A_i, $<_{A_i}$ is a *strict partial order* on the set of worlds $\mathcal{W}$. For the moment, we assume that, in any typicality formula $\mathbf{T}(A)$, A is a distinguished proposition.

Definition 2. *A* (multi-)preferential interpretation *is a triple* $\mathcal{M} = \langle \mathcal{W}, \{<_{A_i}\}, v \rangle$ *where:*

- $\mathcal{W}$ *is a non-empty set of worlds;*
- *each* $<_{A_i} \subseteq \mathcal{W} \times \mathcal{W}$ *is an irreflexive and transitive relation on* $\mathcal{W}$*;*
- $v : \mathcal{W} \longrightarrow 2^{Prop}$ *is a valuation function, assigning to each world* w *a set of propositional variables in* $Prop$.

A *ranked interpretation* is a (multi-)preferential interpretation $\mathcal{M} = \langle \mathcal{W}, \{<_{A_i}\}, v \rangle$ for which all preference relations $<_{A_i}$ are *modular*, that is: for all x, y, z, if $x <_{A_i} y$ then $x <_{A_i} z$ or $z <_{A_i} y$. A relation $<_{A_i}$ is *well-founded* if it does not allow for infinitely descending chains of worlds $w_0, w_1, w_2, \ldots$, with $w_1 <_{A_i} w_0$, $w_2 <_{A_i} w_1$, etc. The valuation v is inductively extended to all formulae:

$$\mathcal{M}, w \models \top \qquad \mathcal{M}, w \not\models \bot$$

$$\mathcal{M}, w \models p \text{ iff } p \in v(w), \quad \text{for all } p \in Prop$$

$$\mathcal{M}, w \models A \wedge B \text{ iff } \mathcal{M}, w \models A \text{ and } \mathcal{M}, w \models B$$

$$\mathcal{M}, w \models A \vee B \text{ iff } \mathcal{M}, w \models A \text{ or } \mathcal{M}, w \models B$$

$$\mathcal{M}, w \models \neg A \text{ iff } \mathcal{M}, w \not\models A$$

$$\mathcal{M}, w \models A \rightarrow B \text{ iff } \mathcal{M}, w \models A \text{ implies } \mathcal{M}, w \models B$$

$$\mathcal{M}, w \models \mathbf{T}(A_i) \text{ iff } \mathcal{M}, w \models A_i \text{ and } \nexists w' \in \mathcal{W} \text{ s.t. } w' <_{A_i} w \text{ and } \mathcal{M}, w' \models A_i.$$

Whether $\mathbf{T}(A_i)$ is satisfied at a world w or not also depends on the other worlds of the interpretation $\mathcal{M}$. Restricting our consideration to modular interpretations, leads to the notions of satisfiability and validity of a formula in the *ranked (or rational)* multi-preferential semantics. Differently from [37] (and from KLM semantics [17, 19]), here we do not assume well-foundedness of the preference relations.

An implication of the form $\mathbf{T}(A) \rightarrow B$, with B in $\mathcal{L}$, corresponds to a conditional $A \mid\sim B$ in KLM logics [17]. It can be easily proven that, when all the preference relations $<_{A_i}$ coincide with a single well-founded preference relation $<$, a multi-preferential interpretation $\mathcal{M}$ corresponds to a KLM preferential interpretation, and a defeasible implication $\mathbf{T}(A) \rightarrow B$ (with A and B in $\mathcal{L}$) has the semantics of a KLM conditional $A \mid\sim B$. The multi-preferential semantics is a generalization of the KLM preferential semantics.

Given a preferential interpretation $\mathcal{M}$, a formula α is *satisfied in* $\mathcal{M}$ if $\mathcal{M}, w \models \alpha$ for some world $w \in \mathcal{W}$. A formula α is *valid in* $\mathcal{M}$ (written $\mathcal{M} \models \alpha$) if $\mathcal{M}, w \models \alpha$, for all the worlds $w \in \mathcal{W}$. A formula α is *valid* if α is valid in all the preferential interpretations $\mathcal{M}$.

Let a *knowledge base* K be a set of (strict or defeasible) implications. A *preferential model of* K is a multi-preferential interpretation $\mathcal{M}$ such that $\mathcal{M} \models A \rightarrow B$, for all implications $A \rightarrow B$ in K. Given a knowledge base K, we say that an implication $A \rightarrow B$ *is preferentially entailed from* K if $\mathcal{M} \models A \rightarrow B$ holds, for all preferential models $\mathcal{M}$ of K. We say that $A \rightarrow B$ *is rationally entailed from* K if $\mathcal{M} \models A \rightarrow B$ holds, for all ranked models $\mathcal{M}$ of K.

It is well known that preferential entailment and rational entailment are weak. As with the rational closure [19] and the lexicographic closure [39] for KML conditionals, also in the multi-preferential case one can strengthen entailment by restricting to specific preferential models, based on some *closure constructions*, which allow to define the preference relations $<_{A_i}$ from a knowledge base K, e.g., by exploiting the ranks or weights of conditional implications, when available [40, 41, 10].

3. A two-valued preferential interpretation of gradual semantics

In [2], we have defined a preferential interpretation of an argumentation graph G under a gradual argumentation semantics S, based on the *many-valued* conditional logic with a multi-preferential semantics. In this section, we construct a conditional interpretation of an argumentation graph G under a gradual semantics S, based on a *two-valued* conditional logic with typicality.

As we have seen, the *semantics* S *of* G can then be regarded, abstractly, as a pair $(\mathcal{D}, \Sigma_G^S)$: a domain of argument valuation $\mathcal{D}$ and a set of labellings Σ_G^S over the domain.

If we consider the set of arguments $\mathcal{A}$ as propositional variables, each labelling σ can be regarded as a world $w_\sigma \in \mathcal{W}$ in a many-valued preferential interpretation $\mathcal{M}$ which contains a preference relation $<_{A_i}$, for each argument A_i (here we are assuming that the single arguments $A_1, \ldots, A_m$ correspond to the distinguished propositions).

More precisely, a gradual semantics $(\mathcal{D}, \Sigma_G)$ of an argumentation graph G can be associated with a preferential interpretation $\mathcal{M}^G = \langle \mathcal{W}, \{<_{A_1}, \ldots, <_{A_n}\}, v \rangle$, defined by letting:

- $\mathcal{W} = \{w_\sigma \mid \sigma \in \Sigma_G\}$
- for all the arguments $A_i \in \mathcal{A}$, and a *threshold value* $t \in \mathcal{D}$:

$$v(w_\sigma, A_i) = \begin{cases} 0 & \textit{if}\, \sigma(A_i) \leq t \\ 1 & \textit{otherwise} \end{cases} \tag{3}$$

- for all the arguments $A_i \in \mathcal{A}$, and worlds $w_\sigma, w_{\sigma'} \in \mathcal{W}$:

$$w_\sigma <_{A_i} w_{\sigma'} \text{ iff } \sigma(A_i) > \sigma'(A_i)$$

The choice of the threshold depends on the semantics and on the domain. For instance, for the domain $\mathcal{D} = [0, 1]$, one may fix the threshold t, e.g., to be 0 or 0.5. For a domain $\mathcal{D} = \mathbb{R}$, one may take $t = 0$. When evaluating the typical worlds (labellings) for A_i, only the worlds in which A_i has a strength above the threshold need to be considered. The preference relation with respect to argument A_i is induced by the strength of the argument A_i in the different labellings.

Example 2. *For instance, referring to the argumentation graph in Example 1, the labelling (strength function)* $\sigma = (3/5, 0, 3/5, 3/5, 1/5, 0)$ *gives rise to a world* w_σ *in the preferential interpretation* $\mathcal{M}^G$ *of the graph. Assuming a threshold* $t = 2/5$*, we have* $v(w_\sigma, A_1) = v(w_\sigma, A_3) = v(w_\sigma, A_4) = 1$ *(as* $\sigma(A_1) > t$*,* $\sigma(A_3) > t$ *and* $\sigma(A_4) > t$*), while* $v(w_\sigma, A_2) = v(w_\sigma, A_5) = v(w_\sigma, A_6) = 0$*.*

Furthermore, for a labelling $\sigma' \in \Sigma_G$ *such that* $\sigma'(A_3) = 4/5$*, we will have:* $w_{\sigma'} <_{A_3} w_\sigma$*, as* $\sigma'(A_3) > \sigma(A_3)$*. That is, labelling* σ' *represents a more typical situation in which argument* A_3 *holds, with respect to labelling* σ*.*

3.1. Conditionals for explanation and boolean combination of arguments

Once a preferential interpretation $\mathcal{M}^G$ of an argumentation graph G with respect to a gradual semantics S, has been constructed, such interpretation can be used in the verification of strict and conditional graded implications (by checking their validity in the model $\mathcal{M}^G$), for explanation, e.g., by validating conditional relations between arguments.

For instance, given a weighted argumentation graph G describing the rules for assigning loans, e.g., involving the arguments $living\_in\_town$, $young$ and $granted\_loan$, and a gradual semantics S, one may want to verify the property

$$\mathbf{T}(granted\_loan) \rightarrow living\_in\_town \wedge young$$

(normally the loan is granted to people living in town and being young) or the property

$$living\_in\_town \wedge young \rightarrow \mathbf{T}(granted\_loan)$$

(living in town and being young implies that normally the loan is granted). The implications above can be checked for validity over the preferential interpretation $\mathcal{M}^G$, constructed from the set of labellings Σ_G of the graph G in the semantics S.

Let us continue the example concerning the argumentation graph in Figure 1, under the φ-coherent argumentation semantics.

Example 3. *As mentioned before, for the weighted argumentation graph in Figure 1, in the* φ*-coherent argumentation semantics there are 36 labellings in case of a domain* $\mathcal{C}_n$ *with* $n = 5$*. Since* A_2 *supports* A_3*, which in turn attacks* A_5*, some relation can be expected between* A_2 *and* A_5*. Assuming a threshold* $t = 2/5$*, the following conditional implication turns out to be valid in the interpretation* $\mathcal{M}_G$*:*

$$\mathbf{T}(A_2) \rightarrow \neg A_5$$

that is, in the situations (labellings) which maximize the acceptability of argument A_2*, argument* $\neg A_5$ *holds. The corresponding strict implication* $A_2 \rightarrow \neg A_5$ *does not hold.*

In this example, we may wonder whether model $\mathcal{M}^G$ *also validates the implication:*

$$\mathbf{T}(A_1 \vee A_2) \rightarrow \neg A_5$$

It turns out, however, that the last conditional implication is outside the language we have defined.

So far we have assumed that the distinguished propositions correspond to single arguments. We lift this condition and allow typicality formulas $\mathbf{T}(A)$, with A a boolean combination of arguments, to include typicality formulas as $\mathbf{T}(A_1 \vee A_2)$ in the example above, and also the typicality formula in the conditional implication:

$$\mathbf{T}(living\_in\_town \wedge young) \rightarrow granted\_loan$$

To deal with these conditionals, we have to extend preferential interpretations by allowing preference relations $<_A$ associated with boolean combinations of arguments A, and generalizing the semantic condition of typicality formulas (in Definition 2) to any boolean combination of arguments A, as follows:

$\mathcal{M}, w \models \mathbf{T}(A)$ iff $\mathcal{M}, w \models A$ and $\nexists w' \in \mathcal{W}$ s.t. $w' <_A w$ and $\mathcal{M}, w' \models A$

For instance, for evaluating $\mathbf{T}(A_1 \vee A_2) \rightarrow \neg A_5$, we need preference relation $<_{A_1 \vee A_2}$.

The definition of of the preference relation $<_A$ for a boolean combination of arguments can rely on the strength of the atomic arguments in the gradual semantics. The strength of boolean arguments in a labelling can be defined inductively from the strength of the atomic arguments A_i in the gradual semantics by exploiting suitable *truth degree functions* $\otimes$, $\oplus$, $\ominus$ in $\mathcal{D}$, as follows:

$$\sigma(\neg A) = \ominus\sigma(A)$$
$$\sigma(A \wedge B) = \sigma(A) \otimes \sigma(B)$$
$$\sigma(A \vee B) = \sigma(A) \oplus \sigma(B)$$

where A and B are boolean combinations of arguments.

When $\mathcal{D}$ is $[0, 1]$ or the finite truth space $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\}$, for an integer $n \geq 1$, $\otimes$, $\oplus$ and $\ominus$ can be chosen as a triangular norm (or *t-norm*), a triangular co-norm (or *s-norm*) and a negation function in some system of many-valued logic [42]. For instance, in the following, for the φ-coherent semantics, we let $a \otimes b = min\{a, b\}$, $a \oplus b = max\{a, b\}$, and $\ominus a = 1 - a$.

Then, the preference relation $<_A$ can be defined as:

$$w_\sigma <_A w_{\sigma'} \text{ iff } \sigma(A) > \sigma'(A)$$

for all worlds $w_\sigma, w_{\sigma'} \in \mathcal{W}$.

We can reconsider the previous example, based on this generalization of the preferential semantics.

Example 4. *Based on the choice of truth degree functions above, one can prove that the conditional implication* $\mathbf{T}(A_1 \vee A_2) \rightarrow \neg A_5$ *is valid in the preferential interpretation of the argumentation graph* G *in Figure 1 under the* φ*-coherent semantics, while the strict implication* $A_1 \vee A_2 \rightarrow \neg A_5$ *is not valid.*

Note that the idea of interpreting the strength function σ as a valuation in a many-valued logic, was previously used in many-valued preferential interpretations [1, 13]. Instead, here it has been exploited in the construction of a two-valued preferential interpretation $\mathcal{M}^G$.

4. A temporal conditional multi-preferential logic

In this section, we extend the two-valued conditional logic in Section 2.3 with the operators of the Linear Time Temporal Logic (LTL) [23].

Compared with the preferential semantics above, the semantics of $LTL^{\mathbf{T}}$ also considers the temporal dimension, through a set of time points in $\mathbb{N}$. The valuation function assigns, at each time point $n \in \mathbb{N}$, a truth value to each propositional variable in a world $w \in \mathcal{W}$; the preference relations $<^n_{A_i}$ (with respect to each distinguished proposition A_i) are relative to time points. Evolution in time may change the valuation of propositions at the worlds, and it may also change the preference relations between worlds (w might represent a typical situation for a student at time point 0, but not at time point 50).

Definition 3. *A* temporal (multi-)preferential interpretation *(or* $LTL^{\mathbf{T}}$ *interpretation) is a triple* $\mathcal{I} = \langle \mathcal{W}, \{<^n_{A_i}\}_{n \in \mathbb{N}}, v \rangle$ *where:*

- $\mathcal{W}$ *is a non-empty set of worlds;*
- *for each* A_i *and* $n \in \mathbb{N}$, $<^n_{A_i} \subseteq \mathcal{W} \times \mathcal{W}$ *is an irreflexive and transitive relation on* $\mathcal{W}$*;*
- $v : \mathbb{N} \times \mathcal{W} \longrightarrow 2^{Prop}$ *is a valuation function assigning, at each time point* n*, a set of propositional variables in* $Prop$ *to each world* $w \in \mathcal{W}$*.*

For $w \in \mathcal{W}$ and $n \in \mathbb{N}$, $v(n, w)$ is the set of the propositional variables which are true in world w at time point n. If there is no $w' \in \mathcal{W}$ s.t. $w' <^n_A w$, we say that w is a *normal situation* for A at time point n.

Given an $LTL^{\mathbf{T}}$ interpretation $\mathcal{I} = \langle \mathcal{W}, \{<^n_{A_i}\}_{n \in \mathbb{N}}, v \rangle$, we define inductively the *truth value of a formula* A *in a world* w *at time point* n (written $\mathcal{I}, n, w \models A$), as follows:

$\mathcal{I}, n, w \models \top \qquad\qquad \mathcal{I}, n, w \not\models \bot$

$\mathcal{I}, n, w \models p$ iff $p \in v(n, w)$, for all $p \in Prop$

$\mathcal{I}, n, w \models A \wedge B$ iff $\mathcal{I}, n, w \models A$ and $\mathcal{I}, n, w \models B$

$\mathcal{I}, n, w \models A \vee B$ iff $\mathcal{I}, n, w \models A$ or $\mathcal{I}, n, w \models B$

$\mathcal{I}, n, w \models \neg A$ iff $\mathcal{I}, n, w \not\models A$

$\mathcal{I}, n, w \models A \rightarrow B$ iff $\mathcal{I}, n, w \models A$ implies $\mathcal{I}, n, w \models B$

$\mathcal{I}, n, w \models XA$ iff $\mathcal{I}, n+1, w \models A$

$\mathcal{I}, n, w \models \Diamond A$ iff there is an $m \geq n$ such that $\mathcal{I}, m, w \models A$

$\mathcal{I}, n, w \models \Box A$ iff for all $m \geq n$, $\mathcal{I}, m, w \models A$

$\mathcal{I}, n, w \models A \mathcal{U} B$ iff there is an $m \geq n$ such that $\mathcal{I}, m, w \models B$ and,
for all k such that $n \leq k < m$, $\mathcal{I}, k, w \models A$

$\mathcal{I}, n, w \models \mathbf{T}(A_i)$ iff $\mathcal{I}, n, w \models A_i$ and $\nexists w' \in \mathcal{W}$ s.t. $w' <^n_{A_i} w$ and $\mathcal{I}, n, w' \models A_i$.

Note that whether a world w represents a typical situation for A_i at a time point n depends on the preference between worlds at time point n.

A *temporal conditional KB* is a set of $LTL^{\mathbf{T}}$ formulas. We evaluate the satisfiability of a temporal graded formula at the initial time point 0 of a temporal preferential interpretation $\mathcal{I}$.

Definition 4. *An $LTL^{\mathbf{T}}$ formula α is* satisfied *in a temporal preferential interpretation $\mathcal{I} = \langle \mathcal{W}, \{<^n_{A_i}\}_{n \in \mathbb{N}}, v \rangle$ if $\mathcal{I}, 0, w \models \alpha$, for some world $w \in \mathcal{W}$. An $LTL^{\mathbf{T}}$ formula α is* valid in a temporal preferential interpretation $\mathcal{I} = \langle \mathcal{W}, \{<^n_{A_i}\}_{n \in \mathbb{N}}, v \rangle$ *if $\mathcal{I}, 0, w \models \alpha$, for all worlds $w \in \mathcal{W}$. An $LTL^{\mathbf{T}}$ formula α is* valid*, if α is valid in all temporal preferential interpretations $\mathcal{I}$. An $LTL^{\mathbf{T}}$ formula α is* satisfiable*, if α is satisfied in some temporal preferential interpretation $\mathcal{I}$.*

It can be shown that the problem of deciding the satisfiability of an $LTL^{\mathbf{T}}$ formula α can be polynomially reduced to the problem of deciding the satisfiability of a concept C_α in the description logic $LTL^{\mathbf{T}}_{\mathcal{ALC}}$ introduced in [28], which extends the temporal description logic $LTL_{\mathcal{ALC}}$ [27] with the typicality operator. $LTL^{\mathbf{T}}_{\mathcal{ALC}}$ has been proven to be decidable when a finite set of well-founded preference relations $<_{A_1}, \ldots, <_{A_m}$ is considered, and concept inclusions are regarded as global temporal constraints. In turn, the decidability of concept satisfiability in $LTL^{\mathbf{T}}_{\mathcal{ALC}}$ relies on the result that concept satisfiability for $LTL_{\mathcal{ALC}}$ w.r.t. TBoxes is in ExpTime (and, actually, it is ExpTime-complete), both with expanding domains [43] and with constant domains [27].

5. Towards a temporal conditional logic for gradual argumentation

As for the non-temporal case, we aim at instantiating the two-valued temporal conditional logic introduced in the previous section to the gradual argumentation setting, to make it suitable for capturing the dynamics of strength functions in time.

In Amgoud and Doder's framework of gradual semantics [8], it is proven that a uniform iterative way of calculating strengths of arguments can be applied to any semantics based on well-defined evaluation methods, for which convergence is guaranteed.

In the following, we consider an iterative formulation of the φ-coherent semantics, and describe a possible construction of a temporal interpretation for it. In particular, starting from an initial score σ_0, one can iteratively define a sequence of labellings $\sigma_0, \sigma_1, \sigma_2, \ldots$, as follows:

$$\sigma_n(A_i) = \begin{cases} \sigma_0(A_i) & \text{if } R^-(A_i) = \emptyset \\ \varphi(W^G_{\sigma_{n-1}}(A_i)) & \text{if } R^-(A_i) \neq \emptyset \end{cases} \tag{4}$$

In the general case, for the φ-coherent semantics one cannot guarantee that the sequence $\sigma_0, \sigma_1, \sigma_2, \ldots$, starting from an initial score σ_0, converges to some φ-coherent labelling, unless the argumentation graph G is acyclic. Convergence conditions for edge-weighted QBAFs have been studied by Potyka, both in the discrete and in the continuous case [11].

Although the sequence of labellings σ_0, σ_1, σ_2, ..., may not converge, one may be interested in verifying temporal properties over the sequence of labellings (or a finite stretch of it). More generally, one may be interested in considering a finite set Σ_0 of possible initial score functions $\{\sigma_0^1, \sigma_0^2, \ldots, \sigma_0^k\}$. The associated sequences σ_0^j, σ_1^j, σ_2^j, ..., one for each initial score function σ_0^j, determine a set of runs, from which a temporal preferential interpretation $\mathcal{I}^G$ of the argumentation graph G can be constructed.

In the sequence of labellings σ_0^j, σ_1^j, σ_2^j, ... obtained from each initial score function $\sigma_0^j \in \Sigma_0$, the labelling σ_n^j is the one obtained from σ_0^j at iteration n.

Let the set of arguments $\mathcal{A}$ be the set of propositional variables of the temporal conditional logic. We can build a temporal preferential interpretation $\mathcal{I}^G = \langle \mathcal{W}, \{<^n_{A_i}\}_{n\in\mathbb{N}}, v\rangle$ of an argumentation graph G, from a set of initial score functions Σ_0, by introducing a world $w^j \in \mathcal{W}$ for each initial score function σ_0^j in Σ_0. The valuation of propositions A_i in a world w^j at time point n, will be determined by the strength $\sigma_n^j(A_i)$ of the argument A_i in the labelling σ_n^j.

More precisely we define a temporal preferential interpretation $\mathcal{I}^G = \langle \mathcal{W}, \{<^n_{A_i}\}_{n\in\mathbb{N}}, v\rangle$ of an argumentation graph G, with respect to Σ_0, by letting:

- $\mathcal{W} = \{w^j \mid \sigma_0^j \in \Sigma_0\}$
- for all the arguments $A_i \in \mathcal{A}$, and a threshold value $t \in \mathcal{D}$:

$$A_i \in v(n, w^j) \text{ iff } \sigma_n^j(A_i) > t;$$

- for all the arguments $A_i \in \mathcal{A}$, and worlds $w^h, w^j \in \mathcal{W}$:

$$w^j <^n_{A_i} w^h \text{ iff } \sigma_n^j(A_i) > \sigma_n^h(A_i).$$

Note that a temporal many-valued interpretation $\mathcal{I}^G = \langle \mathcal{W}, \{<^n_{A_i}\}_{n\in\mathbb{N}}, v\rangle$ can be seen as a sequence of (non-temporal) preferential interpretations $\mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2, \ldots$, where each $\mathcal{M}_n = \langle \mathcal{W}, \{<^n_{A_i}\}, v^n\rangle$ is constructed from all the labelling Σ_n of the argumentation graph G at the iteration n (one for each initial score function in Σ_0), as for $\mathcal{M}^G$ above.

Once a temporal preferential interpretation $\mathcal{I}^G$ has been constructed, the validity of temporal conditional formulas over arguments, such as $\Box(\mathbf{T}(A_1) \rightarrow A_2 \mathcal{U} A_3 \vee A_3)$, can be verified over the constructed preferential interpretation $\mathcal{I}^G$. In the *loan* example, for instance, one may want to check whether normally, young people leaving in town are eventually granted a loan, i.e., $\mathbf{T}(living\_in\_town \wedge young) \rightarrow \Diamond granted\_loan$.

For boolean combinations of arguments in typicality formulas, the non-temporal solution from Section 3.1 extends naturally to the temporal case. Moreover, the approach described here is not limited to the φ-coherent semantics, but can also be adapted to other gradual argumentation semantics with iterative formulations, such as those in [8, 11].

6. Conclusions

In [44], we introduced a many-valued temporal logic with typicality by extending the many-valued conditional logic of [1] with LTL operators. In this paper, we instead develop a two-valued conditional logic with typicality for gradual argumentation, extending it to the temporal case with standard LTL modalities. Compared to the many-valued setting, this framework is conceptually simpler yet still supports expressive reasoning over argumentation graphs. In the non-temporal setting, this approach enables verification of properties of argumentation graphs under different gradual semantics [2]. For the many-valued φ-coherent semantics [32, 13], we have also provided ASP encodings to check graded conditional implications over atomic and boolean argument combinations. Extending these ASP-based techniques to the two-valued temporal case is a natural direction for future work. The temporal extension introduced here enables reasoning about transient properties in the evolution of argument strength, for instance over sequences of labellings from iterative updates. Such reasoning supports model checking or entailment verification of dynamic behaviors, and may aid explainable AI where argumentation evolves over time.

A particularly promising application lies in the analysis of neural networks. Indeed, a multilayer neural network can be interpreted as a weighted knowledge base [9, 10] or as a weighted argumentation graph [12, 11], given the strong semantic relationships also explored in [13]. In the many-valued setting, conditional weighted KBs have already been shown to capture the stationary states of such networks (or suitable approximations thereof) [35, 10, 13], supporting post-hoc symbolic reasoning and verification. The addition of a temporal dimension, as proposed here, opens the way to verifying properties of dynamic behavior — such as how activations evolve, how information propagates, or how explanatory patterns shift over time. From a logical perspective, our approach offers an alternative to other defeasible temporal logics that enrich temporal operators directly, such as those studied in [25, 45]. While those frameworks include tableaux methods and address decidability of fragments with defeasible temporal connectives [24], our logic maintains standard LTL operators and encodes defeasibility entirely through temporal evolution of preferential structures (namely, by allowing preference relations among worlds to change over time). This yields a clean separation between temporal and nonmonotonic aspects, which may facilitate integration with other symbolic reasoning tools. Finally, our work contributes to neuro-symbolic integration [46, 47, 48] by providing a principled framework that links typicality-based conditional logic with neural models in a temporally dynamic and semantically transparent way.

Directions for further research include: developing ASP encodings for the temporal logic proposed here; studying the decidability and complexity of fragments of the logic; investigating convergence and stability in iterative dynamics; applying the framework to *explainability of evolving argumentation graphs* and time-dependent neural-symbolic systems.

Acknowledgments

This research was partially supported by INDAM-GNCS. Mario Alviano was partially supported by the Italian Ministry of University and Research (MUR) under PRIN project PRODE "Probabilistic declarative process mining", CUP H53D23003420006, under PNRR project FAIR "Future AI Research", CUP H23C22000860006, under PNRR project Tech4You "Technologies for climate change adaptation and quality of life improvement", CUP H23C22000370006, and under PNRR project SERICS "SEcurity and RIghts in the CyberSpace", CUP H73C22000880001; by the Italian Ministry of Health (MSAL) under POS projects CAL.HUB.RIA (CUP H53C22000800006) and RADIOAMICA (CUP H53C22000650006); by the Italian Ministry of Enterprises and Made in Italy under project STROKE 5.0 (CUP B29J23000430005); under PN RIC project ASVIN "Assistente Virtuale Intelligente di Negozio" (CUP B29J24000200005); and by the LAIA lab (part of the SILA labs). Mario Alviano is member of Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] 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), volume 3464 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 4–13.

[2] M. Alviano, L. Giordano, D. Theseider Dupré, A many-valued multi-preferential propositional typicality logic and a conditional interpretation for gradual argumentation, J. Logic and Computation (2025). To appear.

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

[4] 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.

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

[6] 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.

[7] 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.

[8] 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.

[9] L. Giordano, From weighted conditionals of multilayer perceptrons to a gradual argumentation semantics, in: 5th Workshop on Advances in Argumentation in Artif. Intell., volume 3086 of *CEUR Workshop Proceedings*, 2021.

[10] M. Alviano, F. Bartoli, M. Botta, R. Esposito, L. Giordano, D. Theseider Dupré, A preferential interpretation of multilayer perceptrons in a conditional logic with typicality, Int. Journal of Approximate Reasoning 164 (2024). URL: https://doi.org/10.1016/j.ijar.2023.109065.

[11] 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.

[12] A. S. d'Avila Garcez, K. Broda, D. M. Gabbay, Symbolic knowledge extraction from trained neural networks: A sound approach, Artif. Intell. 125 (2001) 155–207.

[13] M. Alviano, L. Giordano, D. Theseider Dupré, Weighted knowledge bases with typicality and defeasible reasoning in a gradual argumentation semantics, Intelligenza Artificiale 18 (2024) 153–174.

[14] J. Delgrande, A first-order conditional logic for prototypical properties, Artificial Intelligence 33 (1987) 105–130.

[15] D. Makinson, General theory of cumulative inference, in: Non-Monotonic Reasoning, 2nd International Workshop, Grassau, FRG, June 13-15, 1988, Proceedings, 1988, pp. 1–18.

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

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

[18] 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.

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

[20] 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.

[21] R. Booth, J. B. Paris, A note on the rational closure of knowledge bases with both positive and negative knowledge, Journal of Logic, Language and Information 7 (1998) 165–190. doi:10.1023/A:1008261123028.

[22] G. Kern-Isberner, Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents, volume 2087 of *LNCS*, Springer, 2001.

[23] E. M. Clarke, O. Grumberg, D. A. Peled, Model checking, MIT Press, 1999.

[24] A. Chafik, F. C. Alili, J. Condotta, I. Varzinczak, A one-pass tree-shaped tableau for defeasible LTL, in: TIME 2021, volume 206 of *LIPIcs*, 2021.

[25] A. Chafik, F. C. Alili, J. Condotta, I. Varzinczak, On the decidability of a fragment of preferential LTL, in: TIME 2020, September 23-25, 2020, Bozen-Bolzano, Italy, volume 178 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[26] A. Chafik, Defeasible temporal logics for the specification and verification of exception-tolerant

systems, PhD Thesis, Artois University, 2022.
[27] C. Lutz, F. Wolter, M. Zakharyaschev, Temporal description logics: A survey, in: TIME, 2008, pp. 3–14.
[28] M. Alviano, L. Giordano, D. Theseider Dupré, Preferential temporal description logics with typicality and weighted knowledge bases, in: Proc. 38th Italian Conf. on Computational Logic (CILC 2023), volume 3428 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.
[29] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On rational entailment for propositional typicality logic, Artif. Intell. 277 (2019).
[30] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, ALC+T: a preferential extension of Description Logics, Fundamenta Informaticae 96 (2009) 1–32.
[31] L. Giordano, From weighted conditionals of multilayer perceptrons to a gradual argumentation semantics, in: 5th Workshop on Advances in Argumentation in Artif. Intell., volume 3086 of *CEUR Workshop Proc.*, 2021. Extended version in CoRR abs/2110.03643.
[32] M. Alviano, L. Giordano, D. Theseider Dupré, Preferential reasoning with typicality in ASP over weighted argumentation graphs in a gradual semantics, in: ASPOCP 2023, Proc. ICLP Workshops 2023, CEUR Workshop Proceedings 3437, 2023.
[33] J. Janssen, M. D. Cock, D. Vermeir, Fuzzy argumentation frameworks., in: IPMU 2008, 2008, pp. 513–520.
[34] 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.
[35] 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 *Lecture Notes in Computer Science*, Springer, 2021, pp. 225–242.
[36] D. M. Gabbay, Equational approach to argumentation networks, Argument Comput. 3 (2012) 87–142.
[37] M. Alviano, L. Giordano, D. Theseider Dupré, A framework for conditional reasoning in answer set programming, CoRR abs/2506.03997 (2025). URL: https://doi.org/10.48550/arXiv.2506.03997, accepted for publication at ICLP 2025 as Technical Communication.
[38] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, Preferential Description Logics, in: LPAR 2007, volume 4790 of *LNAI*, Springer, Yerevan, Armenia, 2007, pp. 257–272.
[39] D. J. Lehmann, Another perspective on default reasoning, Ann. Math. Artif. Intell. 15 (1995) 61–82.
[40] L. Giordano, V. Gliozzi, A reconstruction of multipreference closure, Artif. Intell. 290 (2021).
[41] L. Giordano, D. Theseider Dupré, An ASP approach for reasoning in a concept-aware multipreferential lightweight DL, TPLP 10(5) (2020) 751–766.
[42] S. Gottwald, A Treatise on Many-valued Logics, Research Studies Press, 2001.
[43] K. Schild, Combining terminological logics with tense logic, in: EPIA, 1993, pp. 105–120.
[44] M. Alviano, L. Giordano, D. Theseider Dupré, Towards temporal many-valued conditional logics for gradual argumentation: a preliminary report, in: Proc. of the 8th Workshop on Advances in Argumentation in Artificial Intelligence 2024, co-located with AIxIA 2024, volume 3871 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024.
[45] A. Chafik, F. C. Alili, J. Condotta, I. Varzinczak, Defeasible linear temporal logic, J. Appl. Non Class. Logics 33 (2023) 1–51.
[46] L. Serafini, A. S. d'Avila Garcez, Learning and reasoning with logic tensor networks, in: XVth Int. Conf. of the Italian Association for Artificial Intelligence, AI*IA 2016, Genova, Italy, Nov 29 - Dec 1, volume 10037 of *LNCS*, Springer, 2016, pp. 334–348.
[47] L. C. Lamb, A. S. d'Avila Garcez, M. Gori, M. O. R. Prates, P. H. C. Avelar, M. Y. Vardi, Graph neural networks meet neural-symbolic computing: A survey and perspective, in: C. Bessiere (Ed.), Proc. IJCAI 2020, ijcai.org, 2020, pp. 4877–4884.
[48] M. Setzu, R. Guidotti, A. Monreale, F. Turini, D. Pedreschi, F. Giannotti, GlocalX - from local to global explanations of black box AI models, Artif. Intell. 294 (2021) 103457. doi:`10.1016/j.artint.2021.103457`.