Towards Temporal Many-valued Conditional Logics for Gradual Argumentation: a Preliminary Report

Mario Alviano[1] Laura Giordano[2] Daniele Theseider Dupré[2]

1 Università della Calabria, Italy

2 Università del Piemonte Orientale Italy

AI^3 2024: 8th Workshop on Advances in Argumentation in AI

Aims of the talk

- In [NMR 2023] we have proposed a general approach to define a *many-valued preferential interpretation* of *gradual argumentation semantics*.
  - it allows for *conditional reasoning over arguments* and boolean combination of arguments through the *verification of graded (strict or defeasible) implications* over an argumentation graph (with respect to some gradual semantics).
- In this paper we extend the formalism with the temporal operators of Linear Time Temporal Logic (LTL), thus defining a *propositional many-valued temporal logic with typicality*
  - to reason about the *dynamics of a weighted argumentation graph*;
  - to prove properties about the *transient behavior* of a (recurrent) neural network .

The approach

Given an *argumentation graph G* and a *gradual semantics S*, satisfying *weak conditions on the domain* of argument interpretation, we consider:

- a *many-valued propositional logic with typicality*, where arguments play the role of propositional variables (inspired to PTL and DLs with typicality)
- *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" (with α and β *boolean combination of arguments*):

$$\mathbf{T}(granted\_loan) \rightarrow high\_salary \wedge young \geq 0.7$$

- We build a multi-preferential interpretation I_G^S of a graph G under a semantic S
- Verification of conditional properties over I_G^S by *model checking*

Domain of argument interpretation and argumentation graphs: some assumptions

- We let the *domain of argument interpretation* be a set $\mathcal{D}$, equipped with a *preorder relation* $\leq$ [Baroni et al. 2019]
- Let a *(weighted) argumentation graph* be a tuple:

$$G = \langle \mathcal{A}, \mathcal{R}, \sigma_0, \pi \rangle$$

  - $\mathcal{A}$ is a set of *arguments*,
  - $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ a set of *edges*,
  - $\sigma_0 : \mathcal{A} \to \mathcal{D}$ assigns a *base score* of arguments,
  - $\pi : \mathcal{R} \to \mathbb{R}$ is a *weight function* assigning a positive or negative weight to edges.

  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 A when $\pi(B, A)$ is negative.

Many-valued labellings and gradual semantics

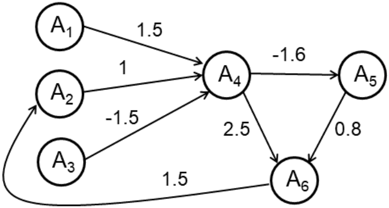

Figure: Example weighted argumentation graph G where the base score is not represented

- A *many-valued labelling (or strength function)* 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* $\mathcal{D}$.
- A *gradual semantics* S for an argumentation graph G identifies *a set* Σ^S *of labellings* of the graph G over a domain of argument valuation $\mathcal{D}$ (considering all possible σ_0).

Example

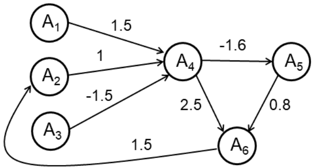

- φ-coherent semantics [NMR 2022]:

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

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

- $\mathcal{D}$ equal to $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, 1\}$.
- With $n = 5$, the graph G has *36 φ-coherent labellings*, while, for $n = 9$, G has *100 φ-coherent 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 $n = 5$.

A many valued logic (of arguments)

- 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* $\mathcal{A}$.
- Language $\mathcal{L}$ contains the boolean connectives $\wedge$, $\vee$, $\neg$ and $\rightarrow$, and that formulas are defined inductively, as usual.
- $\mathcal{D}$ is the *truth degree set*.
- We let $\otimes$, $\oplus$, $\triangleright$ and $\ominus$ be the *truth degree functions* in $\mathcal{D}$ for the connectives $\wedge$, $\vee$, $\neg$ and $\rightarrow$ (respectively).
  - E.g., when $\mathcal{D}$ is $[0, 1]$ or $\mathcal{C}_n$, $\otimes$, $\oplus$, $\triangleright$ and $\ominus$ can be a t-norm, s-norm, implication function, and negation function in some system of many-valued logic.

Many-valued labellings as many-valued valuations

- We can regard a many-valued labelling $\sigma : \mathcal{A} \to \mathcal{D}$ of graph G, assigning to each argument $A_i \in \mathcal{A}$ a truth degree in $\mathcal{D}$, as a *many-valued valuation*.
- σ is extended to all propositional formulas of $\mathcal{L}$:
  $\sigma(\alpha \wedge \beta) = \sigma(\alpha) \otimes \sigma(\beta) \qquad \sigma(\alpha \vee \beta) = \sigma(\alpha) \oplus \sigma(\beta)$
  $\sigma(\alpha \to \beta) = \sigma(\alpha) \rhd \sigma(\beta) \qquad \sigma(\neg\alpha) = \ominus\sigma(\alpha)$
- A labelling σ uniquely assigns a truth degree to any *boolean combination of arguments*.
- We 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.

Preferences over labellings in Σ

- Given a set of labellings Σ, we define a *preference relation* $<_{A_i}$ on Σ, for each argument $A_i \in \mathcal{A}$:

$$\sigma <^{\Sigma}_{A_i} \sigma' \text{ iff } \sigma(A_i) > \sigma'(A_i), \text{ for } \sigma, \sigma' \in \Sigma$$

σ *is* more plausible *than* σ' as a situation for argument A_i to holds.

- The preference relation $<^{\Sigma}_{A_i}$ is a *strict partial order relation* on Σ. We write $<_{A_i}$.
- Similarly, *for boolean combinations of arguments* α:

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

- For example, $\sigma = (1, 4/5, 0, 1, 1/5, 1)$ is preferred to all other labellings with respect to $<_{A_6}$, being the only one with $\sigma(A_6) = 1$.

Preferences with respect to arguments

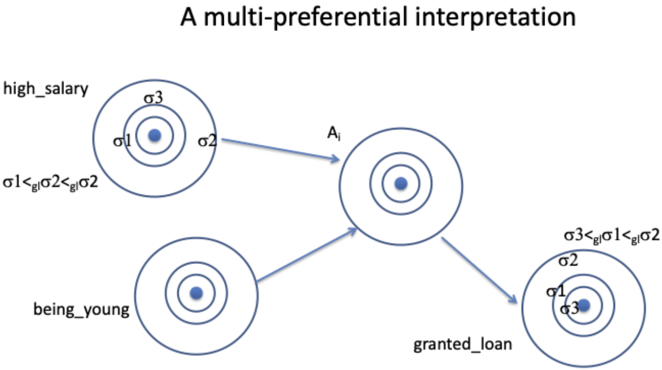

A many-valued logic with typicality

- 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 $I_G^S = (\mathcal{D}, \Sigma^S, \{<_\alpha\})$.
- Language $\mathcal{L}^{\mathbf{T}}$ is obtained by extending $\mathcal{L}$ with a unary *typicality operator* **T**. Intuitively, "a sentence of the form $\mathbf{T}(\alpha)$ is understood to refer to the *typical situations in which* α *holds*" [Booth et al., 2019]
- The typicality operator allows the formulation of *conditional implications* (or *defeasible implications*) of the form $\mathbf{T}(\alpha) \rightarrow \beta$, "normally, if α then β"
- As in PTL also *general implications* $\alpha \rightarrow \beta$, where α and β may contain **T**

A many-valued logic with typicality

- 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:

$$\sigma(\mathbf{T}(\alpha)) = \begin{cases} \sigma(\alpha) & \text{if there is no } \sigma' \text{ such that } \sigma' <_\alpha \sigma \\ 0_{\mathcal{D}} & \text{otherwise} \end{cases} \tag{1}$$

- When $\sigma(\mathbf{T}(A)) > 0_{\mathcal{D}}$, σ is *a labelling maximizing the acceptability of argument A*, among all the labellings in I.

Example

Under Gödel logic with standard involutive negation with $n = 5$, the boolean combination of arguments $A_1 \wedge A_2 \wedge \neg A_3$ has 4 maximally preferred labellings, with $\sigma(A_1 \wedge A_2 \wedge \neg A_3) = 4/5$. For such labellings, $\sigma(\mathbf{T}(A_1 \wedge A_2 \wedge \neg A_3)) = 4/5$, while it is equal to 0 for all other labellings.

Labellings and gradual semantics

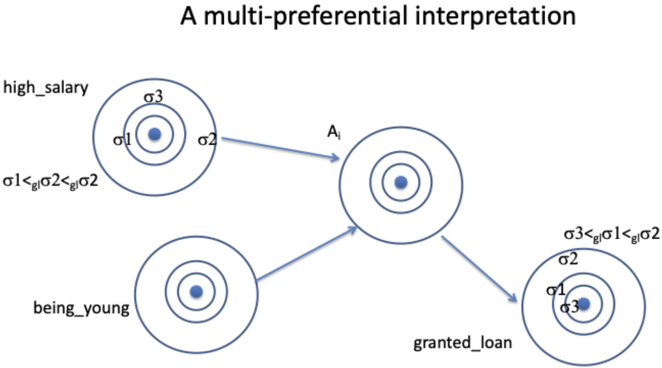

We may check, for instance:

$$\mathbf{T}(\mathit{granted\_loan}) \rightarrow \mathit{high\_salary} \wedge \mathit{being\_young} \geq 0.7$$

Graded implications

- 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.
- the *truth degree of an implication* $\alpha \rightarrow \beta$ *wrt.* I is defined as:

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

- *I satisfies a graded implication* $\alpha \rightarrow \beta \geq t$ (written $I \models \alpha \rightarrow \beta \geq t$) iff $(\alpha \rightarrow \beta)^I \geq t$;

  I satisfies a graded implication $\alpha \rightarrow \beta \leq u$ (written $I \models \alpha \rightarrow \beta \leq u$) iff $(\alpha \rightarrow \beta)^I \leq u$.

Graded implications:example

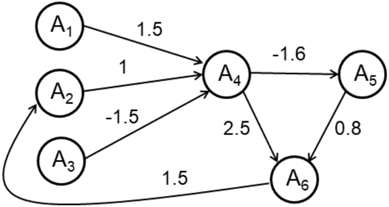

- The following graded conditionals are among the ones satisfied in the preferential interpretation $I = (\mathcal{C}_5, \Sigma, \{<_\alpha\})$, under the φ-coherent semantics:
  $\mathbf{T}(A_1 \wedge A_2 \wedge \neg A_3) \rightarrow A_6 \geq 1$
  (with 4 preferred labellings);
  $\mathbf{T}(A_1 \wedge A_2) \rightarrow A_6 \geq 4/5$ (12 preferred labellings);
  $\mathbf{T}(A_6) \rightarrow A_1 \wedge A_2 \geq 4/5$ (1 preferred labelling).

Properties

Given an interpretation $I^S = (\mathcal{S}, \Sigma^S)$, associated with an argumentation semantics S of a graph G:

- Under the choice of combination functions as in Gödel logic, interpretation $I^S = (\mathcal{S}, \Sigma^S)$ satisfies the *KLM postulates* of a preferential consequence relation, suitably reformulated:

  $\alpha \mid\!\sim \beta$ *is interpreted as* $\mathbf{T}(\alpha) \rightarrow \beta \geq 1$

  $\models A \rightarrow B$ *is interpreted as* $\alpha \rightarrow \beta \geq 1$

- For a *finite* interpretation $I^S = (\mathcal{S}, \Sigma^S)$, *satisfiability* of a graded conditional $\mathbf{T}(\alpha) \rightarrow \beta \geq k$ *in* I^S can be decided in *polynomial time* in the product of the size of the interpretation and the size of the formula.

Temporal multi-preferential interpretations

- We allow temporal operators and the typicality operator to occur in a graded implication. For instance,

$$\mathit{lives\_in\_town} \wedge \mathit{young} \rightarrow \mathbf{T}(\Diamond \mathit{granted\_loan}) \geq 0.8$$
$$\mathbf{T}(\Diamond \mathit{granted\_loan}) \rightarrow \mathit{lives\_in\_town} \wedge \mathit{young} \geq 0.8.$$

- A *temporal (multi-)preferential 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;
  - each $<^n_{A_i} \subseteq \mathcal{W} \times \mathcal{W}$ is partial order on $\mathcal{W}$;
  - $v : \mathbb{N} \times \mathcal{W} \times Prop \longrightarrow \mathcal{D}$ is a valuation function assigning, at each time point, a truth value to any propositional variable (argument) in each world $w \in \mathcal{W}$.

Temporal multi-preferential interpretations

The valuation function v can be extended to all formulas:

$$
\begin{array}{l}
v(n, w, \bot) = 0_{\mathcal{D}} \qquad v(n, w, \top) = 1_{\mathcal{D}} \\
v(n, w, \neg A) = \ominus v(n, w, A) \\
v(n, w, A \wedge B) = v(n, w, A) \otimes v(n, w, B) \\
v(n, w, A \vee B) = v(n, w, A) \oplus v(n, w, B) \\
v(n, w, \mathbf{T}(A)) = \begin{cases} v(n, w, A) & \text{if } \nexists w' \in \mathcal{W} \text{ s.t. } w' <_A^n w \\ 0_{\mathcal{D}} & \text{otherwise} \end{cases} \\
v(n, w, \bigcirc A) = v(n+1, w, A) \\
v(n, w, \Diamond A) = \bigoplus_{m \geq n} v(m, w, A) \\
v(n, w, \Box A) = \bigotimes_{m \geq n} v(m, w, A) \\
v(n, w, A \mathcal{U} B) = \bigoplus_{m \geq n} (v(m, w, B) \otimes \ \bigotimes_{k=n}^{m-1} v(k, w, A))
\end{array}
$$

Following (Frigeri et al. 2014), one can introduce bounded versions for $\Diamond$, $\Box$ and $\mathcal{U}$

Temporal interpretations

We can see a temporal preferential interpretation $\mathcal{I} = \langle \mathcal{W}, \{<^n_{A_i}\}_{n \in \mathbb{N}}, v \rangle$ as a sequence of (non-temporal) preferential interpretations $J^0, J^1, J^2, \ldots$:

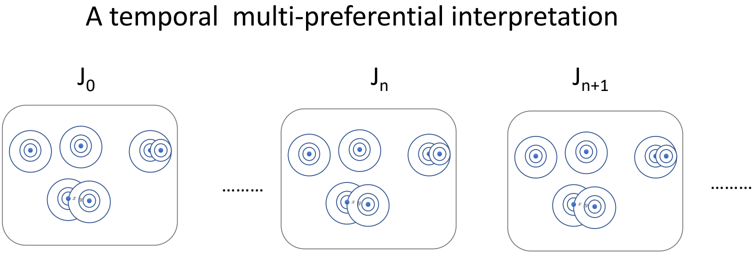

Temporal graded formulas

- Temporal graded implications are evaluated at time point 0:
$$\mathcal{I} \text{ satisfies } A \rightarrow B \geq l \text{ if } (A \rightarrow B)^{\mathcal{I},0} \geq l$$
- Note that either $\mathcal{I}$ satisfies $A \rightarrow B \geq l$ or not: the interpretation above of temporal graded implications in $\mathcal{I}$ at a time point 0 is two-valued.
- Temporal graded formulas can be constructed by combining graded implications:
$$\alpha ::= A \rightarrow B \geq l \mid A \rightarrow B \geq l \mid \alpha \wedge \beta \mid \neg\alpha \mid \\ \bigcirc\alpha \mid \Diamond\alpha \mid \Box\alpha \mid \alpha\mathcal{U}\beta,$$

Example:

$\Box(\mathbf{T}(\mathit{professor}) \rightarrow \mathit{teaches}\,\mathcal{U}\ \mathit{retired} \geq 0.7) \wedge$
$(\mathit{lives\_in\_town} \wedge \mathit{young} \rightarrow \mathbf{T}(\Diamond \mathit{granted\_loan}) \geq 0.8)$

Conclusions and Related work

- We have proposed an approach for *defeasible reasoning* over argumentation graphs in a *temporal formalism*.
- The temporal formalism allows the *dynamics of a weighted argumentation graph* to be captured.
- As a case of study, for the φ-coherent semantics in the finite valued case, the approach has been implemented through an *ASP encoding* [ASPOCP 2023]
- Extending the ASP encodings to deal with temporal preferential interpretations is a direction for future work.