A Conditional and Multi-preferential Approach to Verification and Explanation of Neural Network Models in Answer Set Programming

Mario Alviano[1], Francesco Bartoli[2], Marco Botta[2], Roberto Esposito[2], Laura Giordano[3], Valentina Gliozzi[2] and Daniele Theseider Dupré[3,*]

[1] *Università della Calabria, Italy*
[2] *Università di Torino, Italy*
[3] *Università del Piemonte Orientale, Alessandria, Italy*

Abstract
This short paper reports about a line of research exploiting a conditional logic of commonsense reasoning to provide a semantic interpretation to neural network models. A "concept-wise" multi-preferential semantics for conditionals is exploited to build a preferential interpretation of a trained neural network starting from its input-output behavior. The approach is a general one; it has been first proposed for Self-Organising Maps (SOMs), and then exploited for MultiLayer Perceptrons (MLPs). In fact, a MLP can be regarded as a (fuzzy) conditional knowledge base (KB), in which the synaptic connections correspond to weighted conditionals.

Reasoning for entailment and model-checking in many-valued weighted conditional KBs is based on computational logic (Datalog and Answer Set solving) and it is used in the verification of properties of a trained network and describing what the network has learned.

Keywords
Preferential Description Logics, Typicality, Neural Networks, Explainability, Trustworthiness

1. Introduction

In this short paper we report about our approach to exploit a logic of common sense reasoning for the explainability of some neural network models. We also report about some experiments in the verification of properties of feedforward neural networks by model checking and entailment.

Preferential approaches to common sense reasoning (e.g., [1]) have their roots in conditional logics [2, 3], and have been more recently extended to Description Logics (DLs), to deal with defeasible reasoning in ontologies, by allowing non-strict form of inclusions, called *defeasible* or *typicality* inclusions.

Different preferential semantics [4, 5, 6, 7] and closure constructions (e.g., [8, 9, 10]) have been proposed for defeasible DLs. Among these, the concept-wise multi-preferential semantics [11], which allows to account for preferences with respect to different concepts. It has been introduced first as a semantics of ranked knowledge bases in a lightweight description logic (DL) and then for weighted conditional DL knowledge bases, and proposed as a semantics for some neural network models [12, 13].

We have considered both an unsupervised model, Self-organising maps (SOMs) [14], which is considered a psychologically and biologically plausible neural network model, and a supervised one, MultiLayer Perceptrons (MLPs) [15]. Learning algorithms in the two cases are quite different but our aim was to capture in a semantic interpretation the behavior of the network after training. Considering a domain of input stimuli presented to a network e.g., during training or generalization), a semantic interpretation describing the *input-output behavior* of the network can be provided as a multi-preferential interpretation, where preferences are associated to concepts. For SOMs, the learned categories $C_1, \ldots, C_n$ are regarded as concepts so that a preference relation over the domain of input stimuli is associated with each category [13]. For MLPs, each unit of interest in the deep network (including hidden units) can be associated with a concept and with a preference relation on the domain [12].

For MLPs, the relationship between the logic of commonsense reasoning and deep neural networks is even stronger, as the network can itself be regarded as a conditional knowledge base, i.e., as a set weighted conditionals. This has been achieved by developing a concept-wise *fuzzy multi-preferential semantics* for DLs with weighted defeasible inclusions. Different preferential closure constructions have been considered for weighted knowledge bases (the *coherent* [12], *faithful* [16] and φ-*coherent*

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy
*Corresponding author.
mario.alviano@unical.it (M. Alviano); francesco.bartoli@edu.unito.it (F. Bartoli); marco.botta@unito.it (M. Botta); roberto.esposito@unito.it (R. Esposito); laura.giordano@uniupo.it (L. Giordano); gliozzi@di.unito.it (V. Gliozzi); dtd@uniupo.it (D. Theseider Dupré)

[17] multi-preferential semantics), and their relationships with MLPs have been investigated (see [12, 17]). Undecidability results for fuzzy DLs with general inclusion axioms [18, 19] have motivated the investigation of the (finitely) *many-valued* case. An approach based on Answer Set Programming (ASP) has been proposed for reasoning with weighted conditional KBs under φ-coherent entailment [20]; the complexity of the entailment problem has been studied as well as ASP encodings allowing a solver to deal with KBs corresponding to neural networks with hundreds of nodes and search spaces up to 10^{80} [21]. Datalog with weakly stratified negation has been used for developing a model-checking approach for MLPs in the many-valued case [22]. Both the entailment and the model-checking approaches have been experimented in the verification of properties of trained multilayer feedforward networks.

The strong relationships between neural networks and conditional logics of commonsense reasoning suggest that conditional logics can be used for the verification of properties of neural networks to explain their behavior, in the direction of a trustworthy and explainable AI [23, 24, 25]. The possibility of combining learned knowledge with elicited knowledge in the same formalism is also a step towards neuro-symbolic integration.

2. The concept-wise multi-preferential semantics

The idea underlying the multi-preferential semantics is that, for two domain elements x and y and two concepts, e.g., $Horse$ and $Zebra$, x can be regarded as being more typical than y as a horse ($x <_{Horse} y$), while x could be less typical than y as a zebra ($y <_{Zebra} x$).

This idea has been exploited in the definition of *concept-wise* multi-preferential interpretations [11] for a description logic with typicality concepts (e.g., $\mathbf{T}(Horse)$, representing the class of typical horses), and defeasible inclusions (e.g., $\mathbf{T}(Horse) \sqsubseteq Tall$, meaning that "normally horses are tall"). Typicality inclusions $\mathbf{T}(C) \sqsubseteq D$ correspond to KLM conditionals $C \mathrel{|\!\sim} D$ [1].

Concept-wise multi-preferential interpretations are defined by adding to standard DL interpretations (which are pairs $I = \langle \Delta, \cdot^I \rangle$, where Δ is a domain, and $\cdot^I$ an interpretation function) the preference relations $<_{C_1}$ $, \ldots, <_{C_n}$ associated with a set of distinguished concepts $C_1, \ldots, C_n$, representing the relative typicality of domain individuals with respect to these concepts. Each preference relation $<_{C_i}$ is a modular and well-founded strict partial order on Δ (as preferences in KLM rational models). Preferences with respect to different concepts do not need to agree, as we have seen.

The preference relations are used to define the meaning of typicality concepts. In the two-valued case, a global preference relation $<$ can be defined from the $<_{C_i}$'s, and concept $\mathbf{T}(C)$ is interpreted as the set of all $<$-minimal C elements. In the fuzzy case [12], the preference relation $<_C$ of a concept C is induced by the fuzzy interpretation C^I of the concept, a function mapping each domain element in Δ to a value in $[0, 1]$, that is $x <_C y$ iff $C^I(x) > C^I(y)$.

3. A preferential interpretation of Self-Organising Maps

Once a SOM has learned to categorize, the result of the categorization can be seen as a concept-wise multi-preferential interpretation over a domain of input stimuli, in which a preference relation is associated with each concept (learned category). Once the SOM has learned to categorize, to assess category generalization, Gliozzi and Plunkett [26] define the map's disposition to consider a new stimulus y as a member of a known category C as a function of the *distance* of y from the *map's representation* of C. The distance $d(x, C_i)$ of a stimulus x from a category C_i can be used to build a binary preference relation $<_{C_i}$ among the stimuli in Δ with respect to category C_i [13], by letting $x <_{C_i} y$ if and only if $d(x, C_i) > d(y, C_i)$ (x is more typical than y with respect to category C_i if its distance from category C_i is lower than the relative distance of y). Based on the assumption that the abstraction process in the SOM is able to identify the most typical exemplars for a given category, in the semantic representation of a category, some specific stimuli (corresponding to the *best matching units*) are identified as the *typical exemplars* of the category.

The notion of generalization degree introduced by Gliozzi and Plunkett [26] can be used to define a fuzzy multi-preferential interpretation of SOMs This is done by interpreting each category (concept) as a function mapping each input stimulus to a value in $[0, 1]$, based on the *map's generalization degree* of category membership to the stimulus [26].

In both the two-valued and fuzzy case, the preferential model can be exploited to learn or validate conditional knowledge from empirical data, by verifying conditional formulas over the preferential interpretation constructed from the SOM. In both cases, model checking can be used for the verification of inclusions (either defeasible inclusions or fuzzy inclusion axioms) over the respective models of the SOM (for instance, do the most typical penguins belong to the category Bird with at least a degree of membership 0.8?). Starting from the fuzzy interpretation of the SOM, a probabilistic interpretation of this neural network model is also provided [13], based on Zadeh's probability of fuzzy events [27].

E	F	#counterexamples				#T(E)	P(F/T(E))
		K=1	K=2	K=3	K=4		
Happiness	AU1 ⊔ AU6 ⊔ AU12 ⊔ AU14	0	0	0	22	255	0.8634
	AU6 ⊔ AU12	0	0	1	32	255	0.8422
	AU6 ⊓ AU12	6	15	23	98	255	0.7136
	AU12	0	0	1	35	255	0.8344

Table 1
Results for checking formulae on the test set

4. A preferential interpretation of MultiLayer Perceptrons

The input-output behaviour of MLPs can be captured in a similar way as for SOMs by constructing a preferential interpretation over a domain Δ of input stimuli, e.g., those stimuli considered during training or generalization [12]. Each neuron k of interest for property verification can be associated to a distinguished concept C_k. For each concept C_k, a preference relation $<_{C_k}$ is defined over the domain Δ based on the activity values, $y_k(v)$, of neuron k for each input $v \in \Delta$. In this way, a fuzzy multi-preferential interpretation of the network can be constructed over the domain Δ.

In a fuzzy multi-preferential interpretation, the activation value $y_k(x)$ of neuron k for a stimulus x in the network (assumed to be in the interval $[0, 1]$) is taken to be the degree of membership of x in concept C_k. The interpretation of boolean concepts is defined by fuzzy combination functions, as usual in fuzzy DLs [28, 29]. This also allows a preference relation $<_C$ to be associated to any concept C, and the typical C-elements to be identified, provided the interpretation is well-founded (an assumption which clearly holds when the domain Δ is finite, as in this case). Let us call $\mathcal{M}^{f,\Delta}_{\mathcal{N}}$ the fuzzy multi-preferential interpretation built from network $\mathcal{N}$ over a domain Δ. Logical properties of the network (including fuzzy typicality inclusions) can then be verified by *model checking* over such an interpretation. Evaluating properties involving hidden units might be as well of interest.

A Datalog-based approach has been developed [22], which builds a multi-valued preferential interpretation $\mathcal{M}^{f,\Delta}_{\mathcal{N},n}$ of a trained feedforward network $\mathcal{N}$ and, then, verifies the properties of the network for post-hoc explanation. A multi-valued truth space $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\}$ is considered, for some $n \geq 1$.

The model checking approach has been experimented in the verification of properties of neural networks for the recognition of basic emotions using the Facial Action Coding System (FACS) [30], which involves Action Units (AUs), i.e., facial muscle contractions. From the RAF-DB [31] data set, we selected the subset of the images that were labelled using only one emotion in the set $\{suprise, fear, happiness, anger\}$. A processed dataset containing 5 975 images was input to OpenFace 2.0; the output intensities of AUs were rescaled in order to make their distribution conformant to the expected one in case AUs were recognized by humans [30]. The resulting AUs were used as input to a neural network trained to classify its input as an instance of the four emotions. The neural network model we used is a fully connected feed forward neural network with three hidden layers having 1 800, 1 200, and 600 nodes (all hidden layers use RELU activation functions, while the softmax function is used in the output layer).

The relations between such AUs and emotions, studied by psychologists [32], have been used as a reference for formulae to be verified on neural networks trained to learn such relations. The model checking approach was applied, using the Clingo ASP solver as Datalog engine, taking as set of input stimuli Δ the test set, containing 1194 images, and $n = 5$ (given that AU intensities, when assigned by humans, are on a scale of five values). Table 1 reports some results for the verification of typicality inclusions $\mathbf{T}(E) \sqsubseteq F \geq k/n$, with the number of typical individuals for the emotion E, the number of counterexamples for different values of k (form 1 to n), as well as the value of the conditional probabilities $p(F|\mathbf{T}(E))$ of concept F given concept $\mathbf{T}(E)$, based on Zadeh's probability of fuzzy events [27].

5. MultiLayer Perceptrons as Weighted conditional knowledge bases

The fuzzy multi-preferential interpretation $\mathcal{M}^{f,\Delta}_{\mathcal{N}}$, built from a multilayer perceptron $\mathcal{N}$ for a given set of input stimuli (a domain Δ) as described above, can be proven to be a model of the network $\mathcal{N}$ in a logical sense, by mapping the network into a weighted conditional knowledge base $K^{\mathcal{N}}$ [12].

The weighted conditional knowledge base $K^{\mathcal{N}}$ contains, for each neuron k, a set of weighted defeasible inclusions. If C_k is the concept name associated to unit k and $C_{j_1}, \ldots, C_{j_m}$ are the concept names associated to

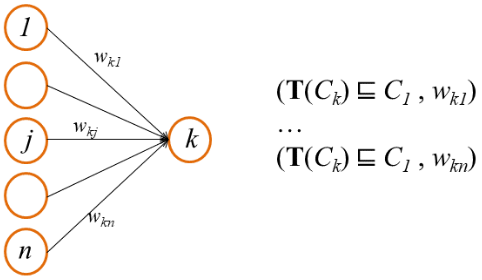

Figure 1: MLPs as knowledge bases

units $j_1, \ldots, j_m$, whose output signals are the input signals for unit k, with synaptic weights $w_{k,j_1}, \ldots, w_{k,j_m}$, then unit k cab be associated a set $\mathcal{T}_{C_k}$ of weighted typicality inclusions: $\mathbf{T}(C_k) \sqsubseteq C_{j_1}$ with $w_{k,j_1}, \ldots$, $\mathbf{T}(C_k) \sqsubseteq C_{j_m}$ with w_{k,j_m} (see figure 1). The fuzzy multipreference interpretation built from a network $\mathcal{N}$ over a domain Δ can be proven to be a model of the knowledge base $K^{\mathcal{N}}$ based on a fuzzy multipreferential semantics, and specifically based on the notions of *coherent* [12], *faithful* [16] and φ*-coherent* [17, 33] (fuzzy) multi-preferential semantics.

In general a weighted conditional KB $K^{\mathcal{N}}$ [12], besides a set of weighted conditional inclusions, also contains a TBox and an ABox as in standard (and in fuzzy) description logics. Multipreferential semantics for weighted conditional KBs have been defined through a semantic closure construction in the spirit of Lehmann's lexicographic closure [34] and Kern-Isberner's c-representations [35], but adopting a concept-wise approach, so that different preference relations are defined.

Specifically, a coherent multi-preferential model of a weighted KB is defined as a fuzzy interpretation $I = \langle \Delta, \cdot^I \rangle$, which satisfies all DL axioms in TBox and ABox, as well as a coherence condition which requires that each preference relation $<_{C_i}$, induced from the fuzzy interpretation over the domain Δ, is coherent with the the weights $W_i(x)$ of all domain individuals x with respect to concept C_i. For each distinguished concept C_i, and domain element $x \in \Delta$, *the weight* $W_i(x)$ *of* x *wrt* C_i in a fuzzy interpretation $I = \langle \Delta, \cdot^I \rangle$ is the sum: $W_i(x) = \sum_h w_h^i \; D_{i,h}^I(x)$.

In the φ*-coherence* semantics a function $\varphi : \mathbb{R} \rightarrow [0, 1]$ is considered (or, more generally, a function φ_i for each distinguished concept C_i). An interpretation $I = \langle \Delta, \cdot^I \rangle$ is φ*-coherent* if, for all concepts $C_i \in \mathcal{C}$ and $x \in \Delta$,

$$C_i^I(x) = \varphi(\sum_h w_h^i \; D_{i,h}^I(x))$$

where $\mathcal{T}_{C_i} = \{(\mathbf{T}(C_i) \sqsubseteq D_{i,h}, w_h^i)\}$ is the set of weighted conditionals for C_i.

If φ is the activation function of a MLP, the activation of neurons for an input vector coincides with the interpretation.

Once a trained neural network can be seen as a weighted defeasible KB $K^{\mathcal{N}}$, φ-entailment can then be used to prove properties of the network for post-hoc explanation.

The model-checking approach does not require to consider the activity of all units, but only of the units involved in the property to be verified; and it only considers a set of input stimuli. Entailment is more challenging from the computational point of view, since all units are considered as well as the space of possible input values. Also in this case, a multi-valued truth space $\mathcal{C}_n = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\}$ is considered, with the approximation φ_n of φ to the closest value in $\mathcal{C}_n$.

The entailment problem has been shown to be $\mathrm{P}^{\mathrm{NP[LOG]}}$-complete [21]. However, encodings in Answer Set Programming have been developed allowing a solver to deal with KBs corresponding to neural networks with hundreds of nodes and search spaces up to 10^{80} (80 input nodes with 10 discrete values) [21].

Some experiments have been done based on finitely many-valued Gödel description logic with typicality $G_n\mathcal{LC}\mathbf{T}$ [20]. As a proof of concept, in [20] the entailment approach has been experimented for the weighted $G_n\mathcal{LC}\mathbf{T}$ KBs corresponding to two of the trained multilayer feedforward network for the MONK's problems ([36]). The approach has also been used for the emotion recognition domain described in section 4, using a smaller network for binary classification of happiness. Using boolean inputs and 10 discrete values for other nodes, the formula

$$\mathbf{T}(happiness) \sqsubseteq au6 \sqcup au12 \; \geq 1$$

has been found to have 4 counterexamples out of the 1446 instances of $\mathbf{T}(happiness)$ among the 2^{17} combinations of boolean inputs. Such 4 combinations do not occur in the data set.

Such cases are interesting as far as trustworthiness is concerned. In cases a formula is expected to hold, we might be satisfied with some exception in the space of input tuples considered for entailment, in a case like the one above, where the exceptional tuples are not represented in the data set. This makes sense in case the data set is a good sample of the real world (as it should be), and the cost of non-compliance with the formula is acceptable. The entailment approach naturally allows for adding constraints on inputs, to exclude part of the input space from the analysis. In other contexts we might be interested in stronger guarantees on previously unseen inputs.

6. Conclusions

Conditional logics of commonsense reasoning can be used for interpreting and verifying the knowledge learned by a neural network for post-hoc explanation and, for MLPs, a trained network can itself be seen as a conditional knowledge base. We see this as a contribution in the direction of a trustworthy and explainable AI [23, 24, 25]. In fact, properties (about input and output concepts) that are found to hold provide a partial description of what the network has learned, i.e., parts of a global explanation [24], where typicality is used to describe properties of cases that are recognized as "strong" members of a class.

Much work has been devoted to the combination of neural networks and symbolic reasoning (e.g., the work by d'Avila Garcez et al. [37, 38, 39] and Setzu et al. [40]), as well as to the definition of new computational models [41, 42, 43, 44]. The work summarized in this paper opens to the possibility of adopting conditional logics as a basis for neuro-symbolic integration, e.g., learning the weights of a conditional knowledge base from empirical data, and combining, for inference, the defeasible inclusions extracted from a neural network with other defeasible or strict inclusions.

References

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

[2] D. Lewis, Counterfactuals, Basil Blackwell Ltd, 1973.

[3] D. Nute, Topics in conditional logic, Reidel, Dordrecht (1980).

[4] 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.

[5] K. Britz, J. Heidema, T. Meyer, Semantic preferential subsumption, in: G. Brewka, J. Lang (Eds.), KR 2008, AAAI Press, Sidney, Australia, 2008, pp. 476–484.

[6] G. Casini, T. A. Meyer, I. Varzinczak, Contextual conditional reasoning, in: AAAI-21, Virtual Event, February 2-9, 2021, AAAI Press, 2021, pp. 6254–6261.

[7] L. Giordano, V. Gliozzi, A reconstruction of multipreference closure, Artif. Intell. 290 (2021).

[8] G. Casini, U. Straccia, Rational Closure for Defeasible Description Logics, in: T. Janhunen, I. Niemelä (Eds.), JELIA 2010, volume 6341 of *LNCS*, Springer, Helsinki, 2010, pp. 77–90.

[9] G. Casini, T. Meyer, K. Moodley, R. Nortje, Relevant closure: A new form of defeasible reasoning for description logics, in: JELIA 2014, LNCS 8761, Springer, 2014, pp. 92–106.

[10] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, Semantic characterization of rational closure: From propositional logic to description logics, Art. Int. 226 (2015) 1–33.

[11] L. Giordano, D. Theseider Dupré, An ASP approach for reasoning in a concept-aware multipreferential lightweight DL, TPLP 10(5) (2020) 751–766.

[12] 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.

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

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

[15] S. Haykin, Neural Networks - A Comprehensive Foundation, Pearson, 1999.

[16] L. Giordano, On the KLM properties of a fuzzy DL with Typicality, in: Proc. ECSQARU 2021, Prague, Sept. 21-24, 2021, volume 12897 of *LNCS*, Springer, 2021, pp. 557–571.

[17] 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. URL: http://ceur-ws.org/Vol-3086/paper8.pdf.

[18] M. Cerami, U. Straccia, On the (un)decidability of fuzzy description logics under Łukasiewicz t-norm, Inf. Sci. 227 (2013) 1–21. URL: https://doi.org/10.1016/j.ins.2012.11.019.

[19] S. Borgwardt, R. Peñaloza, Undecidability of fuzzy description logics, in: G. Brewka, T. Eiter, S. A. McIlraith (Eds.), Proc. KR 2012, Rome, Italy, June 10-14, 2012, AAAI Press, 2012.

[20] 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. doi:10.1017/S1471068422000163.

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

[22] F. Bartoli, M. Botta, R. Esposito, L. Giordano, D. Theseider Dupré, An asp approach for reasoning about the conditional properties of neural networks: an

experiment in the recognition of basic emotions, in: Datalog 2.0 2022: 4th International Workshop on the Resurgence of Datalog in Academia and Industry, volume 3203 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 54–67.
[23] A. Adadi, M. Berrada, Peeking inside the black-box: A survey on explainable artificial intelligence (XAI), IEEE Access 6 (2018) 52138–52160.
[24] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, ACM Comput. Surv. 51 (2019) 93:1–93:42.
[25] A. B. Arrieta, N. D. Rodríguez, J. D. Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, F. Herrera, Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI, Inf. Fusion 58 (2020) 82–115.
[26] V. Gliozzi, K. Plunkett, Grounding bayesian accounts of numerosity and variability effects in a similarity-based framework: the case of self-organising maps, Journal of Cognitive Psychology 31 (2019).
[27] L. Zadeh, Probability measures of fuzzy events, J.Math.Anal.Appl 23 (1968) 421–427.
[28] G. Stoilos, G. B. Stamou, V. Tzouvaras, J. Z. Pan, I. Horrocks, Fuzzy OWL: uncertainty and the semantic web, in: OWLED*05 Workshop on OWL, volume 188 of *CEUR Workshop Proc.*, 2005.
[29] T. Lukasiewicz, U. Straccia, Managing uncertainty and vagueness in description logics for the semantic web, J. Web Semant. 6 (2008) 291–308.
[30] P. Ekman, W. Friesen, J. Hager, Facial Action Coding System, Research Nexus, 2002.
[31] S. Li, W. Deng, J. Du, Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, 2017, pp. 2584–2593.
[32] B. Waller, J. C. Jr., A. Burrows, Selection for universal facial emotion, Emotion 8 (2008) 435–439.
[33] L. Giordano, From weighted conditionals with typicality to a gradual argumentation semantics and back, in: Proc. 20th International Workshop on Non-Monotonic Reasoning, NMR 2022, Part of FLoC 2022, volume 3197 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 127–138.
[34] D. J. Lehmann, Another perspective on default reasoning, Ann. Math. Artif. Intell. 15 (1995) 61–82.
[35] G. Kern-Isberner, Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents, volume 2087 of *LNCS*, Springer, 2001.
[36] Thrun, S. et al., A Performance Comparison of Different Learning Algorithms, Technical Report CMU-CS-91-197, Carnegie Mellon University, 1991.
[37] 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.
[38] A. S. d'Avila Garcez, L. C. Lamb, D. M. Gabbay, Neural-Symbolic Cognitive Reasoning, Cognitive Technologies, Springer, 2009.
[39] A. S. d'Avila Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, S. N. Tran, Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning, FLAP 6 (2019) 611–632.
[40] 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.
[41] 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.
[42] 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.
[43] P. Hohenecker, T. Lukasiewicz, Ontology reasoning with deep neural networks, J. Artif. Intell. Res. 68 (2020) 503–540.
[44] D. Le-Phuoc, T. Eiter, A. Le-Tuan, A scalable reasoning and learning approach for neural-symbolic stream fusion, in: AAAI 2021, February 2-9, AAAI Press, 2021, pp. 4996–5005.