Multilayer Perceptrons as Weighted Conditional Knowledge Bases: an Overview

Laura Giordano[1], Daniele Theseider Dupré[1]

[1] DISIT - Università del Piemonte Orientale, Italy

Abstract
In this paper we report about the relationships between a multi-preferential semantics for defeasible description logics and a deep neural network model. Weighted knowledge bases for description logics are considered under a "concept-wise" preferential semantics, which is further extended to fuzzy interpretations and exploited to provide a preferential interpretation of Multilayer Perceptrons.

Keywords
Common Sense Reasoning, Preferential semantics, Weighted Conditionals, Neural Networks

1. Introduction

Preferential approaches have their roots in conditional logics [1, 2] and have been used to provide axiomatic foundations of non-monotonic and common sense reasoning [3, 4, 5, 6, 7, 8]. More recently they have been extended to description logics (DLs) to deal with inheritance with exceptions in ontologies, by allowing for non-strict forms of inclusions, called *typicality or defeasible inclusions*, with different preferential semantics [9, 10] and closure constructions [11, 12, 13, 14, 15, 16, 17]. This paper exploits a concept-wise multipreference semantics [18] as a semantics for weighted knowledge bases (KBs), i.e. KBs in which defeasible or typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq D$ (meaning "the typical C's are D's" or "normally C's are D's") are given a positive or negative weight.

In this paper we report about the relationships between this logic of common sense reasoning and Multilayer Perceptrons. From the semantic point of view, one can describe the *input-output behavior* of a neural network as a multi-preferential interpretation on the domain of input stimuli, based on the concept-wise multipreference semantics, where preferences are associated to concepts. While in previous work [19, 20], the concept-wise multipreference semantics is used to provide a preferential interpretation of Self-Organising Maps (SOMs) [21], which are regarded as being psychologically and biologically plausible neural network models, in [22] we have investigatesd its relationships with Multilayer Perceptrons (MLPs), a deep neural network model. A deep network is considered after the training phase, when the synaptic weights have been learned, to show that it can be associated a preferential DL interpretation with multiple preferences, as well as a semantics based on fuzzy DL interpretations and another one combining fuzzy interpretations with multiple preferences. The three semantics allow the input-output

AIxIA 2021 Discussion Papers
laura.giordano@uniupo.it (L. Giordano); dtd@uniupo.it (D. Theseider Dupré)

behavior of the network to be captured by interpretations built over a set of input stimuli through a simple construction, which exploits the activity level of neurons for the stimuli. Logical properties can be verified over such models by model checking.

The relationship between the logics of common sense reasoning and Multilayer Perceptrons is even deeper, as a deep neural network can be regarded as a conditional knowledge base with weighted conditionals. This has been achieved by developing a concept-wise fuzzy multipreference semantics for a DL with weighted defeasible inclusions. In the following we recall these results and discuss some challenges from the standpoint of explainable AI [23, 24].

2. A concept-wise multipreference semantics for weighted KBs

A multipreference semantics, taking into account preferences with respect to different concepts, was first introduced by the authors as a semantics for ranked DL knowledge bases [25]. A preference relation $<_{C_i}$ on the domain Δ of a DL interpretation can be associated to each concept C_i to represent the relative typicality of domain individuals with respect to C_i. Preference relations with respect to different concepts do not need to agree, as a domain element x may be more typical than y as a student, but less typical as an employee. The plausibility/implausibility of properties for a concept is represented by their (positive or negative) weight. For instance, a weighted TBox, $\mathcal{T}_{Employee}$, associated to concept $Employee$ might contain the following weighted defeasible inclusions:

(d_1) $\mathbf{T}(Employee) \sqsubseteq Young$, - 50
(d_3) $\mathbf{T}(Employee) \sqsubseteq \exists has\_classes.\top$, -70
(d_2) $\mathbf{T}(Employee) \sqsubseteq \exists has\_boss.Employee$, 100;

meaning that, while an employee normally has a boss, he is not likely to be young or have classes. Furthermore, between the two defeasible inclusions (d_1) and (d_3), the second one is considered to be less plausible than the first one.

Multipreference interpretations are defined by adding to standard DL interpretations, which are pairs $\langle \Delta, \cdot^I \rangle$, 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 $C_1, \ldots, C_n$. Each preference relation $<_{C_i}$ allows for a notion of typicality with respect to concept C_i (e.g. the instances of $\mathbf{T}(Student)$, the typical students, are the preferred domain elements wrt. $<_{Student}$). The definition of a global preference relation $<$ from the $<_{C_i}$'s, leads to the definition of a notion of *concept-wise multipreference interpretation (cwm-interpretation)*, where concept $\mathbf{T}(C)$ is interpreted as the set of all $<$-minimal C-elements. A simple notion of global preference $<$ exploits Pareto combination of the preference relations $<_{C_i}$, but a more sophisticated notion of preference combination has been considered in [18], by taking into account the specificity relation among concepts (e.g., that concept $PhdStudent$ is more specific than concept $Student$). It has been proven [18] that the global preference in a cwm-interpretation determines a KLM-style preferential interpretation, and cwm-entailment satisfies the KLM postulates of a preferential consequence relation [6].

The definition of the concept-wise preferences starting from a weighted conditional knowledge base exploits a closure construction in the same spirit of the one considered by Lehmann [26] to define the lexicographic closure, but more similar to Kern-Isberner's c-representations [27, 28], in which the world ranks are generated as a sum of impacts of falsified conditionals. For weighted $\mathcal{EL}^{\perp}$ knowledge bases [22], the (positive or negative) weights of the satisfied defaults are summed in a concept-wise manner, so to determine the plausibility of a domain elements with respect to certain concepts by considering the modular structure of the KB. Both a two-valued and a fuzzy multipreference semantics have been considered for weighted $\mathcal{EL}^{\perp}$ knowledge bases. In the fuzzy case, to guarantee that the preferences are coherent with the fuzzy interpretation of concepts, a notions of *coherent (fuzzy) multipreference interpretation* has been introduced.

3. A multi-preferential and a fuzzy interpretation for MLPs

Let us consider a deep network after the training phase, when the synaptic weights have been learned. One can describe the *input-output behavior* of the network through a multipreferential interpretation over a (finite) domain Δ of the input stimuli which have been presented to the network during training (or in the generalization phase). The approach is similar to the one proposed for developing a multipreferential interpretation of SOMs [19, 20]. While for SOMs the learned categories are regarded as being DL concepts $C_1, \ldots, C_n$ and each concept C_i is associated a preference relation $<_{C_i}$ over the domain of input stimuli [19, 20] based on a notion of *relative distance* of a stimulus from its *Best Matching Unit* [29], for MLPs, we can associate a concept to each unit of interest, possibly including hidden units. The preference relation associated to a unit is defined based on the activation value of that unit for the different stimuli.

Let $\mathcal{N}$ be a network after training and let $\mathcal{C} = \{C_1, \ldots, C_n\}$ be the set of concept names associated to the units in the network $\mathcal{N}$ we are focusing on. In case the network is not feedforward, we assume that, for each input vector v in Δ, the network reaches a stationary state [30], in which $y_k(v)$ is the activity level of unit k. One can associate to $\mathcal{N}$ and Δ a (two-valued) concept-wise multipreference interpretation over a boolean fragment of $\mathcal{ALC}$ [31] (with no roles and no individual names).

Definition 1. *The* cw^{m}interpretation $\mathcal{M}^{\Delta}_{\mathcal{N}} = \langle \Delta, <_{C_1}, \ldots, <_{C_n}, <, \cdot^I \rangle$ over Δ for network $\mathcal{N}$ *wrt* $\mathcal{C}$ *is a cw^m-interpretation where:*

- *the interpretation function $\cdot^I$ maps each concept name C_k to a set of elements $C_k^I \subseteq \Delta$ and is defined as follows: for all $C_k \in \mathcal{C}$ and $x \in \Delta$, $x \in C_k^I$ if $y_k(x) \neq 0$, and $x \notin C_k^I$ if $y_k(x) = 0$;*
- *for $C_k \in \mathcal{C}$, relation $<_{C_k}$ is defined for $x, x' \in \Delta$ as: $x <_{C_k} x'$ iff $y_k(x) > y_k(x')$, where $y_k(x)$ is the output signal of unit k for input vector x.*

The relation $<_{C_k}$ is a strict modular partial order, and $\leq_{C_k}$ and $\sim_{C_k}$ can be defined as usual. In particular, $x \sim_{C_k} x'$ for $x, x' \notin C_k^I$. Clearly, the boundary between the domain elements which are in C_k^I and those which are not could be defined differently, e.g., by letting $x \in C_k^I$ if $y_k(x) > 0.5$, and $x \notin C_k^I$ if $y_k(x) \leq 0.5$, and suitably adjusting $<_{C_k}$.

This model provides a multipreferential interpretation of the network $\mathcal{N}$, based on the input stimuli considered in Δ, and allows for property verification. For instance, when the neural network is used for categorization and a single output neuron is associated to each category, each concept C_h associated to an output unit h corresponds to a learned category. If

$C_h \in \mathcal{C}$, the preference relation $<_{C_h}$ determines the relative typicality of input stimuli wrt category C_h. This allows to verify typicality properties concerning categories, i.e, $\mathbf{T}(C_h) \sqsubseteq D$ where D is a boolean concept, by *model checking* on the model $\mathcal{M}^{\Delta}_{\mathcal{N}}$. An example is: $\mathbf{T}(\mathit{Eligible\_for\_Loan}) \sqsubseteq \mathit{Lives\_in\_Town} \sqcap \mathit{High\_Salary}$.

Based on the activity level of neurons, a fuzzy DL interpretation can also be constructed. Let N_C be the set of concept names associated to the units of interest in the network $\mathcal{N}$. In a fuzzy DL interpretation $I = \langle \Delta, \cdot^I \rangle$ [32] concepts are interpreted as fuzzy sets over Δ, and the fuzzy interpretation function $\cdot^I$ assigns to each concept $C \in N_C$ a function $C^I : \Delta \rightarrow [0,1]$. For a domain element $x \in \Delta$, $C^I(x)$ represents the degree of membership of x in concept C.

A *fuzzy interpretation* $I_{\mathcal{N}}$ for $\mathcal{N}$ over the domain Δ [22] is a pair $\langle \Delta, \cdot^I \rangle$ where:

(i) Δ is a (finite) set of input stimuli;
(ii) the interpretation function $\cdot^I$ is defined for named concepts $C_k \in N_C$ as: $C_k^I(x) = y_k(x)$, $\forall x \in \Delta$; where $y_k(x)$ is the output signal of neuron k, for input vector x.

The verification that a fuzzy axiom $\langle C \sqsubseteq D \geq \alpha \rangle$ is satisfied in the model $I_{\mathcal{N}}$, can be done based on satisfiability in fuzzy DLs, according to the choice of the fuzzy combination functions. It requires $C_k^I(x)$ to be recorded for all $k = 1, \ldots, n$ and $x \in \Delta$. Of course, one could restrict N_C to the concepts associated to a subset of units, e.g. to input and output units in $\mathcal{N}$ to capture the input/output behavior of the network.

Observe that in a fuzzy interpretation, the interpretation C_h^I of each concept C_h induces an ordering $<_{C_h}$ on the domain Δ, which can be regarded as the preference relation associated to concept C_h. This allows a notion of typicality to be defined in a fuzzy interpretation (in particular, $<_{C_h}$ is well-founded when Δ is finite). The idea underlying fuzzy-multipreference interpretations [22] is to extend a fuzzy DL interpretations with a set of induced preferences, and to identify typical C-elements as the preferred elements wrt. $<_C$. Starting from the fuzzy interpretation of a neural network $\mathcal{N}$, as defined above, a fuzzy-multipreference interpretation $\mathcal{M}^{f,\Delta}_{\mathcal{N}}$ over a domain Δ can be defined, and logical properties of the neural network (combining typicality concepts and fuzzy axioms) can as well be verified over such an interpretations by model checking.

As mentioned in Section 2, fuzzy-multipreference interpretations provide a semantic interpretation of weighted conditional knowledge bases, based on a closure construction. It has been proven that, also in the fuzzy case, the concept-wise multipreference semantics has interesting properties and satisfies most of the KLM properties of a preferential consequence relation, depending of their reformulation in the fuzzy case and on the fuzzy combination functions [33].

The three interpretations considered above for MLPs describe the input-output behavior of the network, and allow for the verification of properties by model-checking. The interpretation $\mathcal{M}^{f,\Delta}_{\mathcal{N}}$ can be proven to be a model of the multilayer network $\mathcal{N}$ when regarded as a weighted conditional KB provided it is *coherent*, i.e., the fuzzy interpretation of concepts agrees with the weights computed from the KB.

Let us assume N_C contains a concept name C_k for each unit k in $\mathcal{N}$. The weighted conditional knowledge base $K^{\mathcal{N}}$ defined from the network $\mathcal{N}$ 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 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 can 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}, ... , $\mathbf{T}(C_k) \sqsubseteq C_{j_m}$ with w_{k,j_m}. The fuzzy multipreference interpretation $\mathcal{M}^{f,\Delta}_{\mathcal{N}}$ built from a network $\mathcal{N}$ and a domain Δ can be proven to be a model of the knowledge base $K^{\mathcal{N}}$ under the some conditions on the activation functions.

4. Conclusions

Much work has been devoted, in recent years, to the combination of neural networks and symbolic reasoning [34, 35, 36], leading to the definition of new computational models [37, 38, 39, 40] and to extensions of logic programming languages with neural predicates [41, 42]. Among the earliest systems combining logical reasoning and neural learning are the KBANN [43] and the CLIP [44] systems and Penalty Logic [45]. The relationships between normal logic programs and connectionist network have been investigated by Garcez and Gabbay [44, 34] and by Hitzler et al. [46]. The correspondence between neural network models and fuzzy systems has been first investigated by Kosko in his seminal work [47]. A fuzzy extension of preferential logics has been studied by Casini and Straccia [48] based on a Rational Closure construction for Gödel fuzzy logic.

The possibility of exploiting the concept-wise multipreference semantics to provide a semantic interpretation of a neural network model has been first explored for Self-Organising Maps, psychologically and biologically plausible neural network models [21]. A multi-preferential semantics can be used to provide a logical model of the SOM behavior after training [19, 20], based on the idea of associating different preference relations to categories, by exploiting the topological organization of the network and a notion of relative distance of an input stimulus from a category. The model can be used to learn or validate conditional knowledge from the empirical data used for training or generalization, by model checking of logical properties. Due to the diversity of the two neural models (MLPs and SOMs), we expect that this approach may be extended to other neural network models and learning approaches.

A logical interpretation of a neural network can be useful from the point of view of explainability, in view of a trustworthy, reliable and explainable AI [23, 24, 49]. For MLPs, the strong relationship between a multilayer network and a weighted KB opens to the possibility of adopting a conditional DLs as a basis for neuro-symbolic integration. While a neural network, once trained, is able and fast in classifying the new stimuli (that is, it is able to do instance checking), all other reasoning services such as satisfiability, entailment and model-checking are missing. These capabilities may be needed to deal with tasks combining empirical and symbolic knowledge, e.g., to extracting knowledge from a network; proving whether the network satisfies (strict or conditional) properties; learning the weights of a conditional KB from empirical data and use them for inference.

To make these tasks possible, the development of proof methods for such logics is a preliminary step. In the two-valued case, multipreference entailment is decidable for weighted $\mathcal{EL}^{\perp}$ KBs [22]. An open problem is whether the notion of fuzzy-multipreference entailment is decidable, for which DLs fragments and under which choice of fuzzy logic combination functions. Undecidability results for fuzzy description logics with general inclusion axioms [50, 51] motivate the investigation of decidable multi-valued approximations of fuzzy-multipreference entailment.

While constructing a conditional interpretation of a neural network is a general approach and can be adapted to different neural network models, it is an issue whether the mapping of deep neural networks to weighted conditional KBs can be extended to more complex neural network models, such as Graph neural networks [37]. Another issue is whether the fuzzy-preferential interpretation of neural networks can be related with the probabilistic interpretation of neural networks based on statistical AI. Indeed, interpreting concepts as fuzzy sets suggests a probabilistic account based on Zadeh's probability of fuzzy events [52], an approach explored by Kosko [47] and exploited for SOMs in [20].

Our work has focused on the multipreference interpretation of MLPs after the learning phase. However, the state of the network during the learning phase can as well be represented as a weighted conditional KB. During training the KB is modified, as weights are updated based on the input stimuli, and one can then regard the learning process as a belief change process. For future work, it would be interesting to study the properties of this notion of change and compare it with the notions of change studied in the literature [53, 54, 55].

Acknowledgments

We thank the anonymous referees for their helpful comments. This research is partially supported by INDAM-GNCS Projects 2020.

References

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

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

[3] D. Gabbay, Theoretical foundations for non-monotonic reasoning in expert systems, Logics and models of concurrent systems, Springer (1985) 439–457.

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

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

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

[7] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial Intelligence 55 (1992) 1–60. doi:http://dx.doi.org/10.1016/0004-3702(92)90041-U.

[8] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: Proc. IJCAI'93, Chambéry, France, August 28 - September 3, 1993, Morgan Kaufmann, ????, pp. 640–647.

[9] 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.

[10] 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.

[11] 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.

[12] G. Casini, T. Meyer, I. J. Varzinczak, , K. Moodley, Nonmonotonic Reasoning in Description Logics: Rational Closure for the ABox, in: 26th International Workshop on Description Logics (DL 2013), volume 1014 of *CEUR Workshop Proceedings*, 2013, pp. 600–615.

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

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

[15] M. Pensel, A. Turhan, Reasoning in the defeasible description logic $EL_{\bot}$ - computing standard inferences under rational and relevant semantics, Int. J. Approx. Reasoning 103 (2018) 28–70.

[16] K. Britz, G. Casini, T. Meyer, K. Moodley, U. Sattler, I. Varzinczak, Principles of KLM-style defeasible description logics, ACM Trans. Comput. Log. 22 (2021) 1:1–1:46.

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

[18] L. Giordano, D. Theseider Dupré, An ASP approach for reasoning in a concept-aware multipreferential lightweight DL, Theory and Practice of Logic programming, TPLP 10(5) (2020) 751–766.

[19] L. Giordano, V. Gliozzi, D. Theseider Dupré, On a plausible concept-wise multipreference semantics and its relations with self-organising maps, in: F. Calimeri, S. Perri, E. Zumpano (Eds.), CILC 2020, Rende, Italy, October 13-15, 2020, volume 2710 of *CEUR*, 2020, pp. 127–140.

[20] L. Giordano, V. Gliozzi, D. Theseider Dupré, A conditional, a fuzzy and a probabilistic interpretation of self-organising maps, CoRR abs/2103.06854 (2021). URL: https://arxiv.org/abs/2103.06854.

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

[22] L. Giordano, D. Theseider Dupré, Weighted defeasible knowledge bases and a multipreference semantics for a deep neural network model, in: Proc17th European Conf. on Logics in AI, JELIA 2021, May 17-20, volume 12678 of *LNCS*, Springer, 2021, pp. 225–242.

[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] L. Giordano, D. Theseider Dupré, An ASP approach for reasoning in a concept-aware multipreferential lightweight DL, Theory Pract. Log. Program. 20 (2020) 751–766.

[26] D. J. Lehmann, Another perspective on default reasoning, Ann. Math. Artif. Intell. 15 (1995) 61–82.

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

[28] G. Kern-Isberner, C. Eichhorn, Structural inference from conditional knowledge bases, Stud Logica 102 (2014) 751–769.

[29] 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).

[30] S. Haykin, Neural Networks - A Comprehensive Foundation, Pearson, 1999.
[31] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider, The Description Logic Handbook - Theory, Implementation, and Applications, 2nd edition, Cambridge, 2007.
[32] T. Lukasiewicz, U. Straccia, Description logic programs under probabilistic uncertainty and fuzzy vagueness, Int. J. Approx. Reason. 50 (2009) 837–853.
[33] L. Giordano, On the KLM properties of a fuzzy DL with Typicality, in: 16th European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU 2021, Springer, 2021. To appear.
[34] 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.
[35] A. S. d'Avila Garcez, L. C. Lamb, D. M. Gabbay, Neural-Symbolic Cognitive Reasoning, Cognitive Technologies, Springer, 2009.
[36] 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.
[37] 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.
[38] 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.
[39] P. Hohenecker, T. Lukasiewicz, Ontology reasoning with deep neural networks, J. Artif. Intell. Res. 68 (2020) 503–540.
[40] 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.
[41] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. D. Raedt, Deepproblog: Neural probabilistic logic programming, in: NeurIPS 2018, 3-8 December 2018, Montréal, Canada, 2018, pp. 3753–3763.
[42] Z. Yang, A. Ishay, J. Lee, Neurasp: Embracing neural networks into answer set programming, in: C. Bessiere (Ed.), Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, ijcai.org, 2020, pp. 1755–1762.
[43] G. G. Towell, J. W. Shavlik, Knowledge-based artificial neural networks, Artif. Intell. 70 (1994) 119–165.
[44] A. S. d'Avila Garcez, G. Zaverucha, The connectionist inductive learning and logic programming system, Appl. Intell. 11 (1999) 59–77.
[45] G. Pinkas, Reasoning, nonmonotonicity and learning in connectionist networks that capture propositional knowledge, Artif. Intell. 77 (1995) 203–247.
[46] P-Hitzler, S. Hölldobler, A. K. Seda, Logic programs and connectionist networks, J. Appl. Log. 2 (2004) 245–272.
[47] B. Kosko, Neural networks and fuzzy systems: a dynamical systems approach to machine intelligence, Prentice Hall, 1992.
[48] G. Casini, U. Straccia, Towards rational closure for fuzzy logic: The case of propositional gödel logic, in: Proc. LPAR-19, Stellenbosch, South Africa, December 14-19, 2013, volume

8312 of *LNCS*, Springer, 2013, pp. 213–227.
[49] 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.
[50] F. Baader, R. Peñaloza, Are fuzzy description logics with general concept inclusion axioms decidable?, in: FUZZ-IEEE 2011, IEEE International Conference on Fuzzy Systems, Taipei, Taiwan, 27-30 June, 2011, Proceedings, IEEE, 2011, pp. 1735–1742.
[51] M. Cerami, U. Straccia, On the undecidability of fuzzy description logics with gcis with lukasiewicz t-norm, CoRR abs/1107.4212 (2011). URL: http://arxiv.org/abs/1107.4212.
[52] L. Zadeh, Probability measures of fuzzy events, J.Math.Anal.Appl 23 (1968) 421–427.
[53] P. Gardenförs, Knowledge in Flux, MIT Press, 1988.
[54] H. Katsuno, A. O. Mendelzon, A unified view of propositional knowledge base updates, in: Proc. IJCAI 1989, Detroit, MI, USA, August 1989, Morgan Kaufmann, 1989, pp. 1413–1419.
[55] H. Katsuno, K. Satoh, A unified view of consequence relation, belief revision and conditional logic, in: IJCAI'91, 1991, pp. 406–412.