

DiSIT, Computer Science Institute
Università del Piemonte Orientale “A. Avogadro”
Viale Teresa Michel 11, 15121 Alessandria
<http://www.di.unipmn.it>



UNIVERSITÀ DEL PIEMONTE ORIENTALE

**Weighted defeasible knowledge bases and a multipreference semantics
for a deep neural network model**

*L. Giordano, D. Theseider Dupré (laura.giordano@uniupo.it,
daniele.theseider_dupre@uniupo.it)*

TECHNICAL REPORT TR-INF-2020-12-05-UNIPMN
(December 2020)

Research Technical Reports published by DiSIT, Computer Science Institute, Università del Piemonte Orientale are available via WWW at URL <http://www.di.unipmn.it/>.
Plain-text abstracts organized by year are available in the directory

Recent Titles from the TR-INF-UNIPMN Technical Report Series

- 2020-04 *A reconstruction of multipreference closure*, L. Giordano, V. Gliozzi, September 2020.
- 2020-03 *A framework for a modular multi-concept lexicographic closure semantics*, L. Giordano, D. Theseider Dupré, September 2020.
- 2020-02 *On a plausible concept-wise multipreference semantics and its relations with self-organising maps*, L. Giordano, V. Gliozzi, D. Theseider Dupré, September 2020.
- 2020-01 *Reasoning about exceptions in ontologies: from the lexicographic closure to the skeptical closure*, L. Giordano, V. Gliozzi, March 2020.
- 2019-05 *Renvoi in Private International Law: a Formalization with Modal Contexts*, L. Giordano, B. Matteo, K. Satoh, October 2019.
- 2019-04 *UML class diagrams supporting formalism definition in the Draw-Net Modeling System*, D. Codetta Raiteri, July 2019.
- 2019-03 *Tracing and preventing sharing and mutation*, P. Giannini, M. Servetto, E. Zucca, July 2019.
- 2019-02 *The Android Forensics Automator (AnForA): a tool for the Automated Forensic Analysis of Android Applications*, C. Anglano, M. Canonico, M. Guazzone, June 2019.
- 2019-01 *Deriving Symbolic and Parametric Structural Relations in Symmetric Nets: Focus on Composition Operator*, L. Capra, M. De Pierro, G. Franceschinis, March 2019.
- 2018-03 *Deriving Symbolic Ordinary Differential Equations from Stochastic Symmetric Nets without Unfolding*, M. Beccuti, L. Capra, M. De Pierro, G. Franceschinis, S. Pernice, July 2018.
- 2018-02 *Power (set) Description Logic*, L. Giordano, A. Policriti, February 2018.
- 2018-01 *A Game-Theoretic Approach to Coalition Formation in Fog Provider Federations (Extended Version)*, C. Anglano, M. Canonico, P. Castagno, M. Guazzone, M. Sereno, February 2018.
- 2017-02 *Configuration and Use of Android Virtual Devices for the Forensic Analysis of Android Applications (see below for citation details)*, C. Anglano, M. Canonico, M. Guazzone, June 2017.
- 2017-01 *A dynamic simulation model for comparing kidney exchange policies*, M. Beccuti, G. Franceschinis, S. Villa, March 2017.
- 2016-04 *Tracing sharing in an imperative pure calculus*, P. Giannini, M. Servetto, E. Zucca, December 2016.
- 2016-03 *SUPPORTING DATA COMMUNICATION AND PATIENT ASSESSMENT DURING EMERGENCY TRANSPORTATION*, M. Canonico, S. Montani, M. Striani, September 2016.

Weighted defeasible knowledge bases and a multipreference semantics for a deep neural network model

Laura Giordano and Daniele Theseider Dupré

DISIT - Università del Piemonte Orientale, Alessandria, Italy

Abstract. In this paper we investigate the relationships between a multipreferential semantics for defeasible reasoning in knowledge representation and a deep neural network model. Weighted knowledge bases for description logics are considered under a “concept-wise” multipreference semantics. The semantics is further extended to fuzzy interpretations and exploited to provide a preferential interpretation of Multilayer Perceptrons, under some condition.

1 Introduction

Preferential approaches have been used to provide axiomatic foundations of non-monotonic and common sense reasoning [18, 40, 43, 35, 44, 36, 5]. 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 [24, 11, 25], and closure constructions [13, 12, 26, 45, 14, 22].

In this paper, we exploit a concept-wise multipreference semantics as a semantics for weighted knowledge bases, i.e. knowledge bases 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. This multipreference semantics, which takes into account preferences with respect to different concepts, has been first introduced as a semantics for ranked DL knowledge bases [20]. For weighted knowledge bases, we develop a different semantic closure construction, although in the spirit of other semantic constructions in the literature. We further extend the multipreference semantics to the fuzzy case.

The concept-wise multipreference semantics has been shown to have some desired properties from the knowledge representation point of view [20, 21], and a related semantics with multiple preferences has also been proposed in the first-order logic setting by Delgrande and Rantsaudis [19]. In previous work [23], the concept-wise multipreference semantics has been used to provide a preferential interpretation of Self-Organising Maps [33], psychologically and biologically plausible neural network models. In this paper, we aim at investigating its relationships with another neural network model, Multilayer Perceptrons.

We consider a multilayer neural network after the training phase, when the synaptic weights have been learned, to show that the neural network can be given a preferential DL semantics 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 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.

To prove that the fuzzy multipreference interpretations, built from the network for a given set of input stimuli, are models of the neural network in a logical sense, we map the multilayer network to a conditional knowledge base, i.e., a set of weighted defeasible inclusions.

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 [1, 27, 2], and can potentially be exploited as the basis for an integrated use of symbolic reasoning and neural models.

2 The description logic \mathcal{ALC} and \mathcal{EL}

In this section we recall the syntax and semantics of the description logic \mathcal{ALC} [4] and of its lightweight fragment \mathcal{EL} [3] at the basis of OWL2 EL Profile.

Let N_C be a set of concept names, N_R a set of role names and N_I a set of individual names. The set of \mathcal{ALC} concepts (or, simply, concepts) can be defined inductively:

- $A \in N_C$, \top and \perp are concepts;
- if C and D are concepts, and $r \in N_R$, then $C \sqcap D$, $C \sqcup D$, $\neg C$, $\forall r.C$, $\exists r.C$ are concepts.

A knowledge base (KB) K is a pair $(\mathcal{T}, \mathcal{A})$, where \mathcal{T} is a TBox and \mathcal{A} is an ABox. The TBox \mathcal{T} is a set of concept inclusions (or subsumptions) $C \sqsubseteq D$, where C, D are concepts. The ABox \mathcal{A} is a set of assertions of the form $C(a)$ and $r(a, b)$ where C is a concept, a an individual name in N_I and r a role name in N_R .

An \mathcal{ALC} interpretation is defined as a pair $I = \langle \Delta, \cdot^I \rangle$ where: Δ is a domain—a set whose elements are denoted by x, y, z, \dots —and \cdot^I is an extension function that maps each concept name $C \in N_C$ to a set $C^I \subseteq \Delta$, each role name $r \in N_R$ to a binary relation $r^I \subseteq \Delta \times \Delta$, and each individual name $a \in N_I$ to an element $a^I \in \Delta$. It is extended to complex concepts as follows:

$$\begin{aligned} \top^I &= \Delta, & \perp^I &= \emptyset, & (\neg C)^I &= \Delta \setminus C^I, \\ (\exists r.C)^I &= \{x \in \Delta \mid \exists y.(x, y) \in r^I \text{ and } y \in C^I\}, & (C \sqcap D)^I &= C^I \cap D^I, \\ (\forall r.C)^I &= \{x \in \Delta \mid \forall y.(x, y) \in r^I \Rightarrow y \in C^I\}, & (C \sqcup D)^I &= C^I \cup D^I. \end{aligned}$$

The notion of satisfiability of a KB in an interpretation and the notion of entailment are defined as follows:

Definition 1 (Satisfiability and entailment). Given an \mathcal{LC} interpretation $I = \langle \Delta, \cdot^I \rangle$:

- I satisfies an inclusion $C \sqsubseteq D$ if $C^I \subseteq D^I$;
- I satisfies an assertion $C(a)$ (resp., $r(a, b)$) if $a^I \in C^I$ (resp., $(a^I, b^I) \in r^I$).

Given a KB $K = (\mathcal{T}, \mathcal{A})$, an interpretation I satisfies \mathcal{T} (resp. \mathcal{A}) if I satisfies all inclusions in \mathcal{T} (resp. all assertions in \mathcal{A}); I is a model of K if I satisfies \mathcal{T} and \mathcal{A} .

A subsumption $F = C \sqsubseteq D$ (resp., an assertion $C(a), r(a, b)$), is entailed by K , written $K \models F$, if for all models $I = \langle \Delta, \cdot^I \rangle$ of K , I satisfies F .

Given a knowledge base K , the *subsumption* problem is the problem of deciding whether an inclusion $C \sqsubseteq D$ is entailed by K .

In the logic \mathcal{EL} [3], concepts are restricted to $C ::= A \mid \top \mid C \sqcap C \mid \exists r.C$, i.e., union, complement and universal restriction are not \mathcal{EL} constructs. In the following, we will also consider the boolean fragment of \mathcal{ALC} only including constructs \sqcap, \sqcup, \neg .

3 Fuzzy description logics

Fuzzy description logics have been widely studied in the literature for representing vagueness in DLs [48, 47, 39, 9, 7], based on the idea that concepts and roles can be interpreted as fuzzy sets and fuzzy binary relations.

As in Mathematical Fuzzy Logic [15] a formula has a degree of truth in an interpretation, rather than being either true or false, in a fuzzy DL axioms are associated with a degree of truth (usually in the interval $[0, 1]$). In the following we shortly recall the semantics of a fuzzy extension of \mathcal{ALC} referring to the survey by Lukasiewicz and Straccia [39]. We limit our consideration to a few features of a fuzzy DL and, in particular, we omit considering datatypes.

A *fuzzy interpretation* for \mathcal{ALC} is a pair $I = \langle \Delta, \cdot^I \rangle$ where: Δ is a non-empty domain and \cdot^I is *fuzzy interpretation function* that assigns to each concept name $A \in N_C$ a function $A^I : \Delta \rightarrow [0, 1]$, to each role name $r \in N_R$ a function $r^I : \Delta \times \Delta \rightarrow [0, 1]$, and to each individual name $a \in N_I$ an element $a^I \in \Delta$. A domain element $x \in \Delta$ belongs to the extension of A to some degree in $[0, 1]$, i.e., A^I is a fuzzy set.

The interpretation function \cdot^I is extended to complex concepts as follows:

$$\begin{aligned} \top^I(x) &= 1, & \perp^I(x) &= 0, & (\neg C)^I(x) &= \ominus C^I(x), \\ (\exists r.C)^I(x) &= \sup_{y \in \Delta} r^I(x, y) \otimes C^I(y), & (C \sqcup D)^I(x) &= C^I(x) \oplus D^I(x) \\ (\forall r.C)^I(x) &= \inf_{y \in \Delta} r^I(x, y) \triangleright C^I(y), & (C \sqcap D)^I(x) &= C^I(x) \otimes D^I(x) \end{aligned}$$

where $x \in \Delta$ and $\otimes, \oplus, \triangleright$ and \ominus are arbitrary but fixed t-norm, s-norm, implication function, and negation function, chosen among the combination functions of various fuzzy logics (we refer to [39] for details).

The interpretation function \cdot^I is also extended to non-fuzzy axioms (i.e., to strict inclusions and assertions of an \mathcal{ALC} knowledge base) as follows:

$$(C \sqsubseteq D)^I = \inf_{x \in \Delta} C^I(x) \triangleright D^I(x), \quad (C(a))^I = C^I(a^I), \quad (R(a, b))^I = R^I(a^I, b^I).$$

A *fuzzy \mathcal{ALC} knowledge base* K is a pair $(\mathcal{T}, \mathcal{A})$ where \mathcal{T} is a fuzzy TBox and \mathcal{A} a fuzzy ABox. A fuzzy TBox is a set of *fuzzy concept inclusions* of the form $C \sqsubseteq D \theta n$, where $C \sqsubseteq D$ is an \mathcal{ALC} concept inclusion axiom, $\theta \in \{\geq, \leq, >, <\}$ and $n \in [0, 1]$. A fuzzy ABox \mathcal{A} is a set of *fuzzy assertions* of the form $C(a)\theta n$ or $r(a, b)\theta n$, where C is an \mathcal{ALC} concept, $r \in N_R$, $a, b \in N_I$, $\theta \in \{\geq, \leq, >, <\}$ and $n \in [0, 1]$. Following Bobillo and Straccia [8], we assume that fuzzy interpretations are *witnessed*, i.e., the sup and inf are attained at some point of the involved domain. The notions of satisfiability of a KB in a fuzzy interpretation and of entailment are defined in the natural way.

Definition 2 (Satisfiability and entailment for fuzzy KBs). A *fuzzy interpretation* I satisfies a fuzzy \mathcal{ALC} axiom E (denoted $I \models E$), as follows, for $\theta \in \{\geq, \leq, >, <\}$:

- I satisfies a fuzzy \mathcal{ALC} inclusion axiom $C \sqsubseteq D \theta n$ if $(C \sqsubseteq D)^I \theta n$;
- I satisfies a fuzzy \mathcal{ALC} assertion $C(a) \theta n$ if $C^I(a^I) \theta n$;
- I satisfies a fuzzy \mathcal{ALC} assertion $r(a, b) \theta n$ if $r^I(a^I, b^I) \theta n$.

Given a fuzzy KB $K = (\mathcal{T}, \mathcal{A})$, a fuzzy interpretation I satisfies \mathcal{T} (resp. \mathcal{A}) if I satisfies all fuzzy inclusions in \mathcal{T} (resp. all fuzzy assertions in \mathcal{A}). A fuzzy interpretation I is a model of K if I satisfies \mathcal{T} and \mathcal{A} . A fuzzy axiom E is entailed by a fuzzy knowledge base K , written $K \models E$, if for all models $I = \langle \Delta, \cdot^I \rangle$ of K , I satisfies E .

4 A concept-wise multipreference semantics for weighted KBs

In this section we develop an extension of \mathcal{EL} with defeasible inclusions having positive and negative weights, based on a concept-wise multipreference semantics first introduced for ranked $\mathcal{EL}_\perp^\dagger$ knowledge bases [20], where defeasible inclusions have positive integer ranks. In addition to standard \mathcal{EL} inclusions $C \sqsubseteq D$ (called *strict* inclusions in the following), the TBox \mathcal{T} will also contain typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq D$, where C and D are \mathcal{EL} concepts. A typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ means that “typical C ’s are D ’s” or “normally C ’s are D ’s” and corresponds to a conditional implication $C \sim D$ in Kraus, Lehmann and Magidor’s (KLM) preferential approach [35, 36]. Such inclusions are defeasible, i.e., admit exceptions, while strict inclusions must be satisfied by all domain elements. We assume that with each typicality inclusion is associated a weight w , a real number. A positive weight supports the plausibility of a defeasible inclusion; a negative weight supports its implausibility.

4.1 Weighted \mathcal{EL} knowledge bases

Let $\mathcal{C} = \{C_1, \dots, C_k\}$ be a set of distinguished \mathcal{EL} concepts, the concepts for which defeasible inclusions are defined. A weighted TBox \mathcal{T}_{C_i} is defined for each distinguished concept $C_i \in \mathcal{C}$ as a set of defeasible inclusions of the form $\mathbf{T}(C_i) \sqsubseteq D$ with a weight.

A *weighted \mathcal{EL} knowledge base K over \mathcal{C}* is a tuple $\langle \mathcal{T}_{strict}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$, where \mathcal{T}_{strict} is a set of strict concept inclusions, \mathcal{A} is an ABox and, for each $C_j \in \mathcal{C}$, \mathcal{T}_{C_j} is a weighted TBox of defeasible inclusions, $\{(d_h^i, w_h^i)\}$, where each d_h^i is a typicality inclusion of the form $\mathbf{T}(C_i) \sqsubseteq D_{i,h}$, having weight w_h^i , a real number.

Consider, for instance, the ranked knowledge base $K = \langle \mathcal{T}_{strict}, \mathcal{T}_{Employee}, \mathcal{T}_{Student}, \mathcal{A} \rangle$, over the set of distinguished concepts $\mathcal{C} = \{Employee, Student\}$, with empty ABox, and with \mathcal{T}_{strict} containing the set of strict inclusions:

$$Employee \sqsubseteq Adult \quad Adult \sqsubseteq \exists has_SSN.\top \quad PhdStudent \sqsubseteq Student$$

The weighted TBox $\mathcal{T}_{Employee}$ contains the following weighted defeasible inclusions:

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

the weighted TBox $\mathcal{T}_{Student}$ contains the defeasible inclusions:

- (d_4) $\mathbf{T}(Student) \sqsubseteq Young, 90$
- (d_5) $\mathbf{T}(Student) \sqsubseteq \exists has_classes.\top, 80$
- (d_6) $\mathbf{T}(Student) \sqsubseteq \exists hasScholarship.\top, -30$

The meaning is 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 less plausible than the first one. Given two employees Tom and Bob such that Tom is not young, has no boss and has classes, while Bob is not young, has a boss and has no classes, in the following, considering the weights above, we will regard Bob as being more typical than Tom as an employee.

4.2 The concept-wise preferences from weighted knowledge bases

The concept-wise multipreference semantics has been recently introduced as a semantics for ranked \mathcal{EL}_\perp^+ knowledge bases [20], which are inspired by Brewka's framework of basic preference descriptions [10]. For each concept $C_i \in \mathcal{C}$, a preference relation $<_{C_i}$ describes the preference among domain elements with respect to concept C_i . Each $<_{C_i}$ has the properties of preference relations in KLM-style ranked interpretations [36], that is, $<_{C_i}$ is a modular and well-founded strict partial order. In particular, $<_{C_i}$ is *well-founded* if, for all $S \subseteq \Delta$, if $S \neq \emptyset$, then $\min_{<_{C_i}}(S) \neq \emptyset$; $<_{C_i}$ is *modular* if, for all $x, y, z \in \Delta$, $x <_{C_j} y$ implies ($x <_{C_j} z$ or $z <_{C_j} y$).

In the following we will recall the concept-wise semantics for \mathcal{ALC} , which extends to its fragments considered in the following. An \mathcal{ALC} interpretation, is extended with a collection of preference relations, one for each concept in \mathcal{C} .

Definition 3 (Multipreference interpretation). A multipreference interpretation is a tuple $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, \cdot^I \rangle$, where:

- (a) Δ is a domain, and \cdot^I an interpretation function, as in \mathcal{ALC} interpretations;
- (b) the $<_{C_i}$ are irreflexive, transitive, well-founded and modular relations over Δ .

The preference relation $<_{C_i}$ determines the relative typicality of domain individuals with respect to concept C_i . For instance, Tom may be more typical than Bob as a student ($tom <_{Student} bob$), but more exceptional as an employee ($bob <_{Employee} tom$). The minimal C_i -elements with respect to $<_{C_i}$ are regarded as the most typical C_i -elements.

While preferences do not need to agree, arbitrary conditional formulas cannot be evaluated with respect to a single preference relation. For instance, evaluating the inclusion “Are typical employed students young?” would require both the preferences $<_{Student}$ and $<_{Employee}$ to be considered. The approach proposed in [20] is that of *combining* the preference relations $<_{C_i}$ into a single *global preference* relation $<$, and then exploit the global preference for interpreting the typicality operator \mathbf{T} , which may be applied to arbitrary concepts. A natural way to define the notion of global preference $<$ is by Pareto combination of the relations $<_{C_1}, \dots, <_{C_k}$, as follows:

$$x < y \text{ iff } \begin{array}{l} (i) \ x <_{C_i} y, \text{ for some } C_i \in \mathcal{C}, \text{ and} \\ (ii) \ \text{for all } C_j \in \mathcal{C}, \ x \leq_{C_j} y. \end{array}$$

A slightly more sophisticated notion of preference combination, which exploits a modified Pareto condition taking into account the specificity relation among concepts (such as, for instance, the fact that concept $PhdStudent$ is more specific than concept $Student$), has been considered for ranked knowledge bases [20].

The addition of the global preference relation, leads to the definition of a notion of *concept-wise multipreference interpretation*, where concept $\mathbf{T}(C)$ is interpreted as the set of all $<$ -minimal C elements.

Definition 4. A concept-wise multipreference interpretation (or cw^m -interpretation) is a multipreference interpretation $\mathcal{M} = \langle \Delta, \langle_{C_1}, \dots, \langle_{C_k}, \langle, \cdot^I \rangle \rangle$, according to Definition 3, such that the global preference relation \langle is defined as above and $(\mathbf{T}(C))^I = \min_{\langle}(C^I)$, where $\text{Min}_{\langle}(S) = \{u : u \in S \text{ and } \nexists z \in S \text{ s.t. } z \langle u\}$.

In the following, we define a notion of cw^m -model of a weighted \mathcal{EL} knowledge base K as a cw^m -interpretation in which the preference relations \langle_{C_i} are constructed from the typicality inclusions in the \mathcal{T}_{C_i} 's.

4.3 A semantics closure construction for weighted knowledge bases

Given a weighted knowledge base $K = \langle \mathcal{T}_{strict}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$, where $\mathcal{T}_{C_i} = \{(d_h^i, w_h^i)\}$ for $i = 1, \dots, k$, and an \mathcal{EL} interpretation $I = \langle \Delta, \cdot^I \rangle$ satisfying all the strict inclusions in \mathcal{T}_{strict} and assertions in \mathcal{A} , we define a preference relation \langle_{C_i} on Δ for each distinguished concepts $C_i \in \mathcal{C}$ through a *semantic closure construction*, a construction similar in spirit to the one considered by Lehmann for the lexicographic closure [37], but based on a different seriousness ordering. In order to define \langle_{C_i} we consider the sum of the weights of the defeasible inclusions for C_i satisfied by each domain element $x \in \Delta$; higher preference wrt \langle_{C_i} is given to the domain elements whose associated sum (wrt C_i) is higher.

First, let us define when a domain element $x \in \Delta$ satisfies/violates a typicality inclusion for C_i wrt an \mathcal{EL} interpretation I . As \mathcal{EL} has the *finite model property* [3], we will restrict to \mathcal{EL} interpretations with a *finite* domain. We say that $x \in \Delta$ *satisfies* $\mathbf{T}(C_i) \sqsubseteq D$ in I , if $x \notin C_i^I$ or $x \in D^I$ (otherwise x *violates* $\mathbf{T}(C_i) \sqsubseteq D$ in I). Note that, in an interpretation I , any domain element which is not an instance of C_i trivially satisfies all defeasible inclusions $\mathbf{T}(C_i) \sqsubseteq D$. Such domain elements will be given the lowest preference with respect to \langle_{C_i} .

Given an \mathcal{EL} interpretation $I = \langle \Delta, \cdot^I \rangle$ and a domain element $x \in \Delta$, we define the *weight of x wrt C_i in I* $W_i(x)$ considering the inclusions $(\mathbf{T}(C_i) \sqsubseteq D_{i,h}, w_h^i) \in \mathcal{T}_{C_i}$:

$$W_i(x) = \begin{cases} \sum_{h: x \in D_{i,h}^I} w_h^i & \text{if } x \in C_i^I \\ -\infty & \text{otherwise} \end{cases} \quad (1)$$

where $-\infty$ is added at the bottom of all real values.

Informally, given an interpretation I , for $x \in C_i^I$, the weight $W_i(x)$ of x wrt C_i is the sum of the weights of all the defeasible inclusions for C_i satisfied by x in I . The more plausible are the satisfied inclusions, the higher is the weight of x . For instance, in the example (Section 4.1), assuming that domain elements $tom, bob \in Employee^I$, and that the typicality inclusion (d_3) is satisfied by tom , while $(d_1), (d_2)$ are satisfied by bob , for $C_i = Employee$, we would get $W_i(tom) = -70$ and $W_i(bob) = 100 - 70 = 30$.

Based on this notion of weight of a domain element wrt a concept, one can construct a preference relation \langle_{C_i} from a given \mathcal{EL} interpretation I . A domain element x is preferred to element y wrt C_i if the weight of the defaults in \mathcal{T}_{C_i} satisfied by x is higher than weight of defaults in \mathcal{T}_{C_i} satisfied by y .

Definition 5 (Preference relation $<_{C_i}$ constructed from \mathcal{T}_{C_i}). Given a ranked knowledge base K where, for all j , $\mathcal{T}_{C_j} = \{(d_h^i, r_h^i)\}$, and an \mathcal{EL} interpretation $I = \langle \Delta, \cdot^I \rangle$, a preference relation \leq_{C_i} can be defined as follows: For $x, y \in \Delta$,

$$x \leq_{C_i} y \text{ iff } W_i(x) \geq W_i(y) \quad (2)$$

\leq_{C_j} is a total preorder relation on Δ . A strict preference relation (a strict modular partial order) $<_{C_j}$ and an equivalence relation \sim_{C_j} can be defined on Δ by letting: $x <_{C_j} y$ iff ($x \leq_{C_j} y$ and not $y \leq_{C_j} x$), and $x \sim_{C_j} y$ iff ($x \leq_{C_j} y$ and $y \leq_{C_j} x$). Note that the domain elements which are instances of C_i are all preferred (wrt $<_{C_i}$) to the domain elements which are not instances of C_i . Furthermore, for all domain elements $x, y \notin C_j^I$, $x \sim_{C_j} y$ holds. The higher is the weight of an element wrt C_i the more preferred is the element. In the example, $W_i(\text{bob}) = 30 > W_i(\text{tom}) = -70$ (for $C_i = \text{Employee}$) and, hence, $\text{bob} <_{\text{Employee}} \text{tom}$, i.e., Bob is more typical than Tom as an employee.

Following the same approach as for ranked \mathcal{EL} knowledge bases [20], we define a notion of cw^m -model for a weighted knowledge base K , where each preference relation $<_{C_j}$ in the model is constructed from the weighted TBox \mathcal{T}_{C_j} according to Definition 5 above, and the global preference is defined by combining the $<_i$'s.

Definition 6 (cw^m -model of K). Let $K = \langle \mathcal{T}_{\text{strict}}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$ be a weighted \mathcal{EL} knowledge base over \mathcal{C} , and $I = \langle \Delta, \cdot^I \rangle$ an \mathcal{EL} interpretation for K . A concept-wise multipreference model (cw^m -model) of K is a cw^m -interpretation $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, <, \cdot^I \rangle$ such that: \mathcal{M} satisfies all strict inclusions in $\mathcal{T}_{\text{strict}}$ and assertions in \mathcal{A} , and for all $j = 1, \dots, k$, $<_{C_j}$ is defined from \mathcal{T}_{C_j} and I , according to Definition 5.

As preference relations $<_{C_j}$, defined according to Definition 5, are irreflexive, transitive, modular, and well-founded relations over Δ (for well-foundedness, remember that we are considering finite models), the notion of cw^m -model \mathcal{M} introduced above is well-defined. By definition of cw^m -model, \mathcal{M} must satisfy all strict inclusions and assertions in K , but it is not required to satisfy all typicality inclusions $\mathbf{T}(C_j) \sqsubseteq D$ in K , unlike other preferential typicality logics [24]. This happens in a similar way in the multipreferential semantics for \mathcal{EL}_\perp^+ ranked knowledge bases, and we refer to [20] for an example showing that the cw^m -semantics is more liberal (in this respect) than standard KLM preferential semantics.

Observe that the notion of weight $W_i(x)$ of x wrt C_i , defined above as the sum of the weights of the satisfied defaults, is just a possible choice for the definition of the preference relations $<_i$ with respect to a concept C_i . A different notion of preference $<_{C_i}$ has been defined from a ranked TBox \mathcal{T}_{C_j} [20], by exploiting the (positive) integer ranks of the defeasible inclusions in \mathcal{T}_{C_j} and the (lexicographic) $\#$ strategy in the framework of basic preference descriptions [10]. The sum of the ranks has been first used in Kern-Isberner's c-interpretations [31, 32], also considering the sum of the weights $\kappa_i^- \in \mathbb{N}$, representing penalty points for falsified conditionals. Here, we only sum the (positive or negative) weights of the satisfied defaults, and we do it in a concept-wise manner.

A notion of *concept-wise entailment* (or cw^m -entailment) can be defined in a natural way to establish when a defeasible concept inclusion follows from a weighted

knowledge base K . We can restrict our consideration to (finite) *canonical* models, i.e., models which are large enough to contain all the relevant domain elements¹.

Definition 7 (cw^m-entailment). An inclusion $\mathbf{T}(C) \sqsubseteq D$ is cw^m-entailed from a weighted knowledge base K if $\mathbf{T}(C) \sqsubseteq D$ is satisfied in all canonical cw^m-models \mathcal{M} of K .

As for ranked \mathcal{EL} knowledge bases [20], it can be proved that this notion of cw^m-entailment for weighted KBs satisfies the KLM postulates of a preferential consequence relation [20]. This is an easy consequence of the fact that the global preference relation $<$, which is used to evaluate typicality, is a strict partial order. As $<$ is not necessarily modular, cw^m-entailment does not necessarily satisfy rational monotonicity [36].

The problem of deciding cw^m-entailment is Π_2^P -complete for ranked \mathcal{EL}_\perp^+ knowledge bases [20]; cw^m-entailment can be proven as well to be in Π_2^P for weighted knowledge bases, based on a similar reformulation of cw^m-entailment as a problem of computing preferred answer sets. The proof of the result is similar to the proof of Proposition 7 in the online Appendix of [20], apart from minor differences due to the different notion of preference $<_{C_i}$ used here with respect to the one for ranked knowledge bases.

5 Weighted Tboxes and multipreference fuzzy interpretations

In this section, we move to consider fuzzy interpretations, and investigate the possibility of extending the previous multipreference semantic construction to the fuzzy case.

Definition 8 (Fuzzy multipreference interpretation). A fuzzy multipreference interpretation (or fm-interpretation) is a tuple $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, \cdot^I \rangle$, where:

- (a) (Δ, \cdot^I) is a fuzzy interpretation;
- (b) the $<_{C_i}$ are irreflexive, transitive, well-founded and modular relations over Δ ;

Let K be a weighted knowledge base $\langle \mathcal{T}_{strict}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$, where each axiom in \mathcal{T}_{strict} has the form $\langle \alpha \geq 1 \rangle$, and $\mathcal{T}_{C_i} = \{(d_h^i, w_h^i)\}$ is a set of typicality inclusions $d_h^i = \mathbf{T}(C_i) \sqsubseteq D_{i,h}$ with weight w_h^i .

Given a fuzzy interpretation $I = \langle \Delta, \cdot^I \rangle$, satisfying all the strict inclusions in \mathcal{T}_{strict} and all assertions in \mathcal{A} , we aim at constructing a concept-wise multipreference interpretation from I , by defining a preference relation $<_{C_j}$ on Δ for each $C_i \in \mathcal{C}$, based on a closure construction similar to the one developed in Section 4.3. The definition of $W_i(x)$ in (1) can be reformulated as follows:

$$W_i(x) = \begin{cases} \sum_h w_h^i D_{i,h}^I(x) & \text{if } C_i^I(x) > 0 \\ -\infty & \text{otherwise} \end{cases} \quad (3)$$

by regarding the interpretation $D_{i,h}^I$ of concept $D_{i,h}$ as a two valued function from Δ to $\{0, 1\}$ (rather than a subset of Δ). And similarly for $C_i^I(x)$. Definition (3) can be taken as the definition of the weight function $W_i(x)$ when I is a fuzzy interpretation.

¹ This is a standard assumption in the semantic characterizations of rational closure for DLs, and in other semantic constructions. See [20] for the definition of canonical models for \mathcal{EL} .

Simply, in the fuzzy case, for each default $d_h^i = \mathbf{T}(C_i) \sqsubseteq D_{i,h}$, $D_{i,h}^I(x)$ is a value in $[0, 1]$. In the sum, the value $D_{i,h}^I(x)$ of the membership of x in $D_{i,h}$ is weighted by w_h^i . For inclusions $\mathbf{T}(C_i) \sqsubseteq D_{i,h}$ with a positive weight, the higher is the degree of truth of the membership of x in $D_{i,h}$, the higher is the weight $W_i(x)$. For inclusions with a negative weight, the lower is the degree of truth of the membership of x in $D_{i,h}$, the higher is the weight $W_i(x)$.

From this notion of weight of a domain element x wrt a concept $C_i \in \mathcal{C}$, the *preference relation* \leq_{C_i} associated with \mathcal{T}_{C_j} in a fuzzy interpretation I can be defined as in Section 4.3:

$$x \leq_{C_i} y \text{ iff } W_i(x) \geq W_i(y) \quad (4)$$

A notion of fuzzy multipreference model of a weighted KB can then be defined.

Definition 9 (fuzzy multipreference model of K). Let $K = \langle \mathcal{T}_{strict}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$ be a weighted \mathcal{EL} knowledge base over \mathcal{C} . A fuzzy multipreference model (or fm-model) of K is an fm-interpretation $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, \cdot^I \rangle$ such that: the fuzzy interpretation $I = \langle \Delta, \cdot^I \rangle$ satisfies all strict inclusions in \mathcal{T}_{strict} and assertions in \mathcal{A} and, for all $j = 1, \dots, k$, $<_{C_j}$ is defined from \mathcal{T}_{C_j} and I , according to condition (4).

Note that, as we restrict to witnessed fuzzy interpretations I , for $S \neq \emptyset$, $\min_{x \in S} C_i^I$ is attained at some point in Δ . Hence, $\min_{<_{C_i}}(S) \neq \emptyset$, i.e., $<_{C_i}$ is well-founded.

The preference relation $<_{C_i}$ establishes how typical a domain element x is wrt C_i . We can then require that the degree of membership in C_i (given by the fuzzy interpretation I) and the relative typicality wrt C_i (given by the preference relations $<_{C_i}$) are related, and agree with each other.

Definition 10 (Coherent fm-models). The preference relation $<_{C_i}$ agrees with the fuzzy interpretation $I = \langle \Delta, \cdot^I \rangle$ if, for all $x, y \in \Delta$: $x <_{C_i} y$ iff $C_i^I(x) > C_i^I(y)$. An fm-model $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, \cdot^I \rangle$ of K is a coherent fm-model (or cf^m -model) of K if, for all $C_i \in \mathcal{C}$, preference relation $<_{C_i}$ agrees with the fuzzy interpretation I .

In a cf^m -model, the preference relation $<_{C_i}$ over Δ constructed from \mathcal{T}_{C_i} coincides with the preference relation induced by C_i^I . As the interpretation function \cdot^I extends to any concept C , for cf^m -models we do not need to introduce a global preference relation $<$, defined by combining the $<_{C_i}$. To define the interpretation of typicality concepts $\mathbf{T}(C)$ in a cf^m -model, we follow a different route and we let, for all concepts C ,

$$(\mathbf{T}(C))^I = \min_{<_C}(C^I),$$

where $<_C$ is the preference relation over Δ induced by C^I , i.e., for all $x, y \in \Delta$: $x <_C y$ iff $C^I(x) > C^I(y)$. Note that $\mathbf{T}(C)$ is a two valued concept, i.e., $(\mathbf{T}(C))^I(x) \in \{0, 1\}$, and satisfiability in a cf^m -model is now extended to fuzzy inclusion axioms involving typicality concepts, such as $\langle \mathbf{T}(C) \sqsubseteq D \geq \alpha \rangle$.

A notion of cf^m -entailment from a weighted knowledge base K can be defined in the obvious way: a fuzzy axiom E is cf^m -entailed by a fuzzy knowledge base K if, for all cf^m -models \mathcal{M} of K , \mathcal{M} satisfies E .

6 Preferential and fuzzy interpretations of multilayer perceptrons

In this section, we first shortly introduce multilayer perceptrons. Then we develop a preferential interpretation of a neural network after training, and a fuzzy preferential interpretation.

Let us first recall from [28] the model of a *neuron* as an information-processing unit in an (artificial) neural network. The basic elements are the following:

- a set of *synapses* or *connecting links*, each one characterized by a *weight*. We let x_j be the signal at the input of synapse j connected to neuron k , and w_{kj} the related synaptic weight;
- the adder for summing the input signals to the neuron, weighted by the respective synapses weights: $\sum_{j=1}^n w_{kj}x_j$;
- an *activation function* for limiting the amplitude of the output of the neuron (typically, to the interval $[0, 1]$ or $[-1, +1]$).

The sigmoid, threshold and hyperbolic-tangent functions are examples of activation functions. A neuron k can be described by the following pair of equations: $u_k = \sum_{j=1}^n w_{kj}x_j$, and $y_k = \varphi(u_k + b_k)$, where x_1, \dots, x_n are the input signals and w_{k1}, \dots, w_{kn} are the weights of neuron k ; b_k is the bias, φ the activation function, and y_k is the output signal of neuron k . By adding a new synapse with input $x_0 = +1$ and synaptic weight $w_{k0} = b_k$, one can write:

$$u_k = \sum_{j=0}^n w_{kj}x_j \quad y_k = \varphi(u_k), \quad (5)$$

where u_k is called the *induced local field* of the neuron. The neuron can be represented as a directed graph, where the input signals x_1, \dots, x_n and the output signal y_k of neuron k are nodes of the graph. An edge from x_j to y_k , labelled w_{kj} , means that x_j is an input signal of neuron k with synaptic weight w_{kj} .

A neural network can then be seen as “a directed graph consisting of nodes with interconnecting synaptic and activation links” [28]: nodes in the graph are the neurons (the processing units) and the weight w_{ij} on the edge from node j to node i represents “the strength of the connection [...] by which unit j transmits information to unit i ” [42]. Source nodes (i.e., nodes without incoming edges) produce the input signals to the graph. Neural network models are classified by their synaptic connection topology. In a *feedforward* network the architectural graph is acyclic, while in a *recurrent* network it contains cycles. In a feedforward network neurons are organized in layers. In a *single-layer* network there is an input-layer of source nodes and an output-layer of computation nodes. In a *multilayer feedforward* network there is one or more hidden layer, whose computation nodes are called *hidden neurons* (or hidden units). The source nodes in the input-layer supply the activation pattern (*input vector*) providing the input signals for the first layer computation units. In turn, the output signals of first layer computation units provide the input signals for the second layer computation units, and so on, up to the final output layer of the network, which provides the overall response of the network to the activation pattern. In a recurrent network at least one feedback exists, so that “the

output of a node in the system influences in part the input applied to that particular element” [28]. In the following, we do not put restrictions on the topology the network.

“A major task for a neural network is to learn a model of the world” [28]. In supervised learning, a set of input/output pairs, input signals and corresponding desired response, referred as training data, or training sample, is used to train the network to learn. In particular, the network learns by changing the synaptic weights, through the exposition to the training samples. After the training phase, in the generalization phase, the network is tested with data not seen before. “Thus the neural network not only provides the implicit model of the environment in which it is embedded, but also performs the information-processing function of interest” [28]. In the next section, we try to make this model explicit as a multipreference model.

6.1 A multipreference interpretation of multilayer perceptrons

Assume that the network \mathcal{N} has been trained and the synaptic weights w_{kj} have been learned. We associate a concept name $C_i \in N_C$ to any unit i in \mathcal{N} (including input units and hidden units) and construct a multi-preference interpretation over a (finite) domain Δ of input stimuli, the input vectors considered so far, for training and generalization. In case the network is not feedforward, we assume that, for each input vector v in Δ , the network reaches a stationary state [28], in which $y_k(v)$ is the activity level of unit k .

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a subset of N_C , the set of concepts C_i for a distinguished subset of units i , the units we are focusing on (for instance, \mathcal{C} might be associated to the set of output units, or to all units). We can associate to \mathcal{N} and Δ a (two-valued) concept-wise multipreference interpretation over the boolean fragment of \mathcal{ALC} (with no roles and no individual names), based on Definition 4, as follows:

Definition 11. *The cw^m interpretation $\mathcal{M}_{\mathcal{N}}^{\Delta} = \langle \Delta, \langle_{C_1}, \dots, \langle_{C_n}, \langle, \cdot^I \rangle$ over Δ for network \mathcal{N} wrt \mathcal{C} is a cw^m -interpretation where:*

- *the interpretation function \cdot^I is defined for named concepts $C_k \in N_C$ as: $x \in C_k^I$ if $y_k(x) \neq 0$, and $x \notin C_k^I$ if $y_k(x) = 0$.*
- *for $C_k \in \mathcal{C}$, relation \langle_{C_k} is defined for $x, x' \in \Delta$ as: $x \langle_{C_k} x'$ iff $y_k(x) > y_k(x')^2$.*

The relation \langle_{C_k} is a strict partial order, and \leq_{C_k} and \sim_{C_k} are 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 $y_k(x) \leq 0.5$. This would require only a minor change in the definition of the \langle_{C_k} .

This model provides a multipreference interpretation of the network \mathcal{N} , based on the input stimuli considered in Δ . 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 \langle_{C_h} determines the relative typicality of input stimuli wrt category C_i . This allows to verify typicality properties concerning categories, such as $\mathbf{T}(C_h) \sqsubseteq D$ (where D is a boolean concept built from the named concepts in N_C), by *model*

² $y_k(x)$ is the output signal of unit k for input vectors x . Differently from condition (5), here (and below) the dependency of the output y_k of neuron k on the input vector x is made explicit.

checking on the model $\mathcal{M}_{\mathcal{N}}^{\Delta}$. According to the semantics of typicality concepts, this would require to identify typical C_h -elements and checking whether they are instances of concept D . General typicality inclusion of the form $\mathbf{T}(C) \sqsubseteq D$, with C and D boolean concepts, can as well be verified on the model $\mathcal{M}_{\mathcal{N}}^{\Delta}$. However, the identification of $<$ -minimal C -elements requires computing, for all pairs of elements $x, y \in \Delta$, the relation $<$ and the relations $<_{C_i}$ for $C_i \in \mathcal{C}$. This may be challenging as Δ can be large.

Evaluating properties involving hidden units might be of interest, although their meaning is usually unknown. In the well known Hinton's family example [29], one may want to verify whether, normally, given an old Person 1 and relationship Husband, Person 2 would also be old, i.e., $\mathbf{T}(Old_1 \sqcap Husband) \sqsubseteq Old_2$ is satisfied. Here, concept Old_1 (resp., Old_2) is associated to a (known, in this case) hidden unit for Person 1 (and Person 2), while Husband is associated to an input unit.

6.2 A fuzzy interpretation of multilayer perceptrons

The definition of a fuzzy model of a neural network \mathcal{N} , under the same assumptions as in Section 6.1, is straightforward. Let N_C be the set containing a concept name C_i for each unit i in \mathcal{N} , including hidden units. Let us restrict to the boolean fragment of \mathcal{ALC} with no individual names. We define a *fuzzy interpretation* $I_{\mathcal{N}} = \langle \Delta, \cdot^I \rangle$ for \mathcal{N} as follows:

- Δ is a (finite) set of input stimuli;
- 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 t-norm and implication function. It requires $C_k^I(x)$ to be recorded for all $k = 1, \dots, n$ and $x \in \Delta$. Of course, one could restrict N_C to the concepts associated to input and output units in \mathcal{N} , so to capture the input/output behavior of the network.

In the next section, starting from this fuzzy interpretation of a neural network \mathcal{N} , we define a fuzzy multipreference interpretation $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$, and prove that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a coherent fm-model of the conditional knowledge base $K_{\mathcal{N}}$ associated to \mathcal{N} , under some condition.

6.3 Multilayer perceptrons as conditional knowledge bases

Let N_C be as in Section 6.2, and let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a subset of N_C . Given the *fuzzy interpretation* $I_{\mathcal{N}} = \langle \Delta, \cdot^I \rangle$ as defined in Section 6.2, a fuzzy multipreference interpretation $\mathcal{M}_{\mathcal{N}}^{f,\Delta} = \langle \Delta, <_{C_1}, \dots, <_{C_n}, \cdot^I \rangle$ over \mathcal{C} can be defined by letting $<_{C_k}$ to be the preference relation induced by the interpretation $I_{\mathcal{N}}$, as follows: for $x, x' \in \Delta$,

$$x <_{C_k} x' \text{ iff } y_k(x) > y_k(x'). \quad (6)$$

Interpretation $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ makes the preference relations induced by $I_{\mathcal{N}}$ explicit. We aim at proving that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is indeed a coherent fm-model of the neural network \mathcal{N} . A weighted conditional knowledge base $K^{\mathcal{N}}$ is associated to the neural network \mathcal{N} as follows.

For each unit k , we consider all the units j_1, \dots, j_m whose output signals are the input signals of unit k , with synaptic weights $w_{k,j_1}, \dots, w_{k,j_m}$. Let C_k be the concept name associated to unit k and C_{j_1}, \dots, C_{j_m} the concept names associated to units j_1, \dots, j_m , respectively. We define for each unit k the following set \mathcal{T}_{C_k} of typicality inclusions, with their associated weights:

$$\mathbf{T}(C_k) \sqsubseteq C_{j_1} \text{ with } w_{k,j_1}, \dots, \mathbf{T}(C_k) \sqsubseteq C_{j_m} \text{ with } w_{k,j_m}$$

Given \mathcal{C} , the knowledge base extracted from network \mathcal{N} is defined as the tuple: $K^{\mathcal{N}} = \langle \mathcal{T}_{strict}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_n}, \mathcal{A} \rangle$, where $\mathcal{T}_{strict} = \mathcal{A} = \emptyset$ and $K^{\mathcal{N}}$ contains the set \mathcal{T}_{C_k} of weighted typicality inclusions associated to neuron k (defined as above), for each $C_k \in \mathcal{C}$. $K^{\mathcal{N}}$ is a weighted knowledge base over the set of distinguished concepts $\mathcal{C} = \{C_1, \dots, C_n\}$. For multilayer feedforward networks, $K^{\mathcal{N}}$ correspond to an acyclic conditional knowledge base, and defines a (defeasible) subsumption hierarchy among concepts. It can be proved that:

Proposition 1. $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a cf^m -model of the knowledge base $K^{\mathcal{N}}$, provided the activation functions φ of all neurons are monotonically increasing and have value in $(0, 1]$.

Proof. Let $\mathcal{M}_{\mathcal{N}}^{f,\Delta} = \langle \Delta, <_{C_1}, \dots, <_{C_n}, <, \cdot^I \rangle$. Let us consider any neuron k , such that $C_k \in \mathcal{C}$.

From the hypothesis, for any input stimulus $x \in \Delta$, $y_k(x) > 0$ and, hence, $C_k(x) > 0$. The weight $W_k(x)$ of x wrt C_k is defined, according to equation (3), as

$$W_k(x) = \sum_{h=1}^m w_{k,j_h} C_{j_h}^I(x)$$

where $\mathbf{T}(C_k) \sqsubseteq C_{j_1} \dots \mathbf{T}(C_k) \sqsubseteq C_{j_m}$ are the typicality inclusions in \mathcal{T}_{C_k} with weights $w_{k,j_1}, \dots, w_{k,j_m}$. Observe that, for all $h = 1, \dots, m$, $C_{j_h}^I(x) = y_{j_h}$, for input x , by construction of interpretation $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ (and of \cdot^I). Therefore:

$$W_k(x) = \sum_{h=1}^m w_{k,j_h} y_{j_h}$$

As y_{j_1}, \dots, y_{j_m} are the input signals of unit k , it holds that $W_k(x)$ is the induced local field u_k of unit k (as in equation (5)), for the input stimulus x . As $u_k = W_k(x)$ then, for the given input stimulus x , the output of neuron k must be $y_k(x) = \varphi(u_k) = \varphi(W_k(x))$, where φ is the activation function of unit k . As by construction of $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$, $C_k^I(x) = y_k(x)$ for all units k , it holds that $C_k^I(x) = \varphi(W_k(x))$.

To prove that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a cf^m -model of $K^{\mathcal{N}}$, we have to prove that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a fuzzy multipreference model of $K_{\mathcal{N}}$ and it is coherent.

We first prove that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is an fm-model of $K_{\mathcal{N}}$. As \mathcal{T}_{strict} and \mathcal{A} are empty in $K^{\mathcal{N}}$. We only have to prove that, for all $C_k \in \mathcal{C}$, $<_{C_k}$ satisfies Condition (4). We prove that,

$$x <_{C_k} x' \text{ iff } W_k(x) > W_k(x'),$$

from which (4) follows. For all $x, x' \in \Delta$, by construction of $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$,

$$x <_{C_k} x' \text{ iff } y_k(x) > y_k(x')$$

Assume that $x <_{C_k} x'$. As $y_k(x) = \varphi(W_k(x))$ and $y_k(x') = \varphi(W_k(x'))$, $\varphi(W_k(x)) > \varphi(W_k(x'))$ holds. Then, it must be the case that $W_k(x) > W_k(x')$ (otherwise, by the assumption that φ is monotone increasing, from $W_k(x) \leq W_k(x')$ it would follow that $\varphi(W_k(x)) \leq \varphi(W_k(x'))$).

Conversely, assume that $W_k(x) > W_k(x')$. As, from the hypothesis, $C_k(x') = y_k(x') > 0$, both $W_k(x)$ and $W_k(x')$ are weighted sum of real numbers. As φ is monotonically increasing, $\varphi(W_k(x)) > \varphi(W_k(x'))$ and, hence, $y_k(x) > y_k(x')$, so that $x <_{C_k} x'$ holds.

It is easy to prove that $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is coherent. By construction of $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$, for each $C_k \in \mathcal{C}$, $<_{C_k}$ is defined by equation (6) as:

$$x <_{C_k} x' \text{ iff } y_k(x) > y_k(x')$$

As $C_k^I(x) = y_k(x)$ and $C_k^I(x') = y_k(x')$ (again by construction of $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$),

$$x <_{C_k} x' \text{ iff } C_k^I(x) > C_k^I(x')$$

i.e., $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a coherent fm-model of $K^{\mathcal{N}}$. □

Under the given conditions, that hold, for instance, for the sigmoid activation function, for any choice of $\mathcal{C} \subseteq N_C$ and for any choice of the domain Δ of input stimuli (all leading to a stationary state of \mathcal{N}), the fm-interpretation $\mathcal{M}_{\mathcal{N}}^{f,\Delta}$ is a coherent fuzzy multipreference model of the defeasible knowledge base $K^{\mathcal{N}}$. The knowledge base $K^{\mathcal{N}}$ does not provide a logical characterization of the neural network \mathcal{N} , as the requirement of coherence does not determine the activation functions of neurons. For this reason, the knowledge base $K^{\mathcal{N}}$ captures the behavior of all the networks \mathcal{N}' , obtained from \mathcal{N} by replacing the activation function of the units in \mathcal{N} with other monotonically increasing activation functions with values in $(0, 1]$, in all possible ways (but retaining the same synaptic weights as in \mathcal{N}). That is, an interpretation $\mathcal{M}_{\mathcal{N}'}^{f,\Delta}$, constructed from a network \mathcal{N}' and any Δ as above, is as well a cf^m -model of $K^{\mathcal{N}}$. This means that the logical formulas cf^m -entailed from $K^{\mathcal{N}}$ hold in all the models $\mathcal{M}_{\mathcal{N}'}^{f,\Delta}$ built from \mathcal{N}' . They are properties of \mathcal{N}' , as well as of network \mathcal{N} . cf^m -entailment from $K^{\mathcal{N}}$ is sound for \mathcal{N} and for each \mathcal{N}' as above.

7 Conclusions

In this paper, we have investigated the relationships between defeasible knowledge bases, under a fuzzy multipreference semantics, and multilayer neural networks. Given a network after training, we have seen that one can construct a (fuzzy) multipreference interpretation starting from a domain containing a set of input stimuli, and using the activity level of neurons for the stimuli. We have proven that such interpretations are models of the conditional knowledge base associated to the network, corresponding to a set of weighted defeasible inclusions in a simple DL.

The correspondence between neural network models and fuzzy systems has been first investigated by Bart Kosko in his seminal work [34]. In his view, “at each instant the

n-vector of neuronal outputs defines a fuzzy unit or a fit vector. Each fit value indicates the degree to which the neuron or element belongs to the n-dimensional fuzzy set.” As a difference, our fuzzy interpretation of a multilayer perceptron regards each concept (representing a single neuron) as a fuzzy set. This is the usual way of viewing concepts in fuzzy DLs [48, 38, 6], and we have used fuzzy concepts within a multipreference semantics based on a semantic closure construction, in the line of Lehmann’s semantics for lexicographic closure [37] and of Kern-Isberner’s c-interpretations [31, 32].

Much work has been devoted, in recent years, to the combination of neural networks and symbolic reasoning, leading to the definition of new computational models [17, 16, 46, 30], and to extensions of logic programming languages with neural predicates [41, 49]. Here, rather than developing a new neural model to capture symbolic reasoning, we have provided a multipreference semantics for a neural model as it is. This logical interpretation may be of interest from the standpoint of explainable AI [1, 27, 2].

Several issues may deserve further investigation as future work. An open problem is whether the notion of c^m -entailment is decidable (even for the small fragment of \mathcal{EL} without roles), under which choice of fuzzy logic combination functions, and whether decidable approximations can be defined. Another issue is whether the multipreference semantics can provide a semantic interpretation of other neural network models, besides self-organising maps [33], whose multipreference semantics has been investigated in [23].

References

1. A. Adadi and M. Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6:52138–52160, 2018.
2. A. Barredo Arrieta, N. Díaz Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera. Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion*, 58:82–115, 2020.
3. F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In L.P. Kaelbling and A. Saffiotti, editors, *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005)*, pages 364–369, Edinburgh, Scotland, UK, August 2005. Professional Book Center.
4. F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, and P.F. Patel-Schneider. *The Description Logic Handbook - Theory, Implementation, and Applications, 2nd edition*. Cambridge, 2007.
5. S. Benferhat, D. Dubois, and H. Prade. Possibilistic logic: From nonmonotonicity to logic programming. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty, European Conference, ECSQARU’93, Granada, Spain, November 8-10, 1993, Proceedings*, pages 17–24, 1993.
6. F. Bobillo and U. Straccia. The fuzzy ontology reasoner fuzzydl. *Knowl. Based Syst.*, 95:12–34, 2016.
7. F. Bobillo and U. Straccia. Reasoning within fuzzy OWL 2 EL revisited. *Fuzzy Sets Syst.*, 351:1–40, 2018.
8. Fernando Bobillo and Umberto Straccia. Reasoning within fuzzy OWL 2 EL revisited. *Fuzzy Sets Syst.*, 351:1–40, 2018.
9. S. Borgwardt, F. Distel, and R. Peñaloza. The limits of decidability in fuzzy description logics with general concept inclusions. *Artif. Intell.*, 218:23–55, 2015.

10. G. Brewka. A rank based description language for qualitative preferences. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004, Valencia, Spain, August 22-27, 2004*, pages 303–307, 2004.
11. K. Britz, J. Heidema, and T. Meyer. Semantic preferential subsumption. In G. Brewka and J. Lang, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 11th International Conference (KR 2008)*, pages 476–484, Sidney, Australia, September 2008. AAAI Press.
12. G. Casini, T. Meyer, I. J. Varzinczak, , and 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*, pages 600–615, 2013.
13. G. Casini and U. Straccia. Rational Closure for Defeasible Description Logics. In T. Janhunen and I. Niemelä, editors, *Proc. 12th European Conf. on Logics in Artificial Intelligence (JELIA 2010)*, volume 6341 of *LNCS*, pages 77–90, Helsinki, Finland, September 2010. Springer.
14. G. Casini, U. Straccia, and T. Meyer. A polynomial time subsumption algorithm for nominal safe elo_{\perp} under rational closure. *Inf. Sci.*, 501:588–620, 2019.
15. P. Cintula, P. Hájek, and C. Noguera, editors. *Handbook of Mathematical Fuzzy Logic*, volume 37-38. College Publications, 2011.
16. A. S. d’Avila Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, and Son N. Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *FLAP*, 6(4):611–632, 2019.
17. A. S. d’Avila Garcez, L. C. Lamb, and D. M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Cognitive Technologies. Springer, 2009.
18. J. Delgrande. A first-order conditional logic for prototypical properties. *Artificial Intelligence*, 33(1):105–130, 1987.
19. J. Delgrande and C. Rantsoudis. A preference-based approach for representing defaults in first-order logic. In *Proc. 18th Int. Workshop on Non-Monotonic Reasoning, NMR2020, September 12th - 14th, 2020*.
20. L. Giordano and D. Theseider Dupré. An ASP approach for reasoning in a concept-aware multipreferential lightweight DL. *Theory Pract. Log. Program.*, 20(5):751–766, 2020. Online Appendix in CoRR, abs/2006.04387, 2020.
21. L. Giordano and D. Theseider Dupré. A framework for a modular multi-concept lexicographic closure semantics. volume abs/2009.00964, 2020. 18th International Workshop on Non-Monotonic Reasoning, NMR2020, September 12th-14th 2020.
22. L. Giordano and V. Gliozzi. L. Giordano and L. Gliozzi: A reconstruction of the multipreference closure. . *Artificial Intelligence*, 290, January 2021. to appear.
23. L. Giordano, V. Gliozzi, and D. Theseider Dupré. On a plausible concept-wise multipreference semantics and its relations with self-organising maps. *CoRR*, abs/2008.13278, 2020. CILC (Italian Conference on Computational Logic), 13-15 October 2020, Rende.
24. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Preferential Description Logics. In Nachum Dershowitz and Andrei Voronkov, editors, *Proceedings of LPAR 2007 (14th Conference on Logic for Programming, Artificial Intelligence, and Reasoning)*, volume 4790 of *LNAI*, pages 257–272, Yerevan, Armenia, October 2007. Springer-Verlag.
25. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. ALC+T: a preferential extension of Description Logics. *Fundamenta Informaticae*, 96:1–32, 2009.
26. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence*, 226:1–33, 2015.
27. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, 2019.

28. S. Haykin. *Neural Networks - A Comprehensive Foundation*. Pearson, 1999.
29. G. Hinton. Learning distributed representation of concepts. In *Proceedings 8th Annual Conference of the Cognitive Science Society*. Erlbaum, Hillsdale, NJ, 1986.
30. P. Hohenecker and T. Lukasiewicz. Ontology reasoning with deep neural networks. *J. Artif. Intell. Res.*, 68:503–540, 2020.
31. G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Springer, 2001.
32. G. Kern-Isberner and C. Eichhorn. Structural inference from conditional knowledge bases. *Stud Logica*, 102(4):751–769, 2014.
33. T. Kohonen, M.R. Schroeder, and T.S. Huang, editors. *Self-Organizing Maps, Third Edition*. Springer Series in Information Sciences. Springer, 2001.
34. Bart Kosko. *Neural networks and fuzzy systems: a dynamical systems approach to machine intelligence*. Prentice Hall, 1992.
35. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.
36. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.
37. D. J. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
38. T. Lukasiewicz and U. Straccia. Managing uncertainty and vagueness in description logics for the semantic web. *J. Web Semant.*, 6(4):291–308, 2008.
39. T. Lukasiewicz and U. Straccia. Description logic programs under probabilistic uncertainty and fuzzy vagueness. *Int. J. Approx. Reason.*, 50(6):837–853, 2009.
40. David Makinson. General theory of cumulative inference. In *Non-Monotonic Reasoning, 2nd International Workshop, Grassau, FRG, June 13-15, 1988, Proceedings*, pages 1–18, 1988.
41. R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 3753–3763, 2018.
42. P. McLeod, K. Plunkett, and E.T. Rolls, editors. *Introduction to Connectionist Modelling of Cognitive Processes*. Oxford university Press, 1998.
43. J. Pearl. *Probabilistic Reasoning in Intelligent Systems Networks of Plausible Inference*. Morgan Kaufmann, 1988.
44. J. Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK'90), Pacific Grove, CA, USA, March 1990*, pages 121–135. Morgan Kaufmann, 1990.
45. M. Pensel and A. Turhan. Reasoning in the defeasible description logic EL_{\perp} - computing standard inferences under rational and relevant semantics. *Int. J. Approx. Reasoning*, 103:28–70, 2018.
46. L. Serafini and A. S. d'Avila Garcez. Learning and reasoning with logic tensor networks. In *AI*IA 2016: Advances in Artificial Intelligence - XVth Int. Conf. of the Italian Association for Artificial Intelligence, Genova, Italy, November 29 - December 1, 2016, Proceedings*, volume 10037 of *LNCS*, pages 334–348. Springer, 2016.
47. G. Stoilos, G. B. Stamou, V. Tzouvaras, J. Z. Pan, and I. Horrocks. Fuzzy OWL: uncertainty and the semantic web. In *Proceedings of the OWLED*05 Workshop on OWL: Experiences and Directions, Galway, Ireland, November 11-12, 2005*, volume 188 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2005.

48. U. Straccia. Towards a fuzzy description logic for the semantic web (preliminary report). In *The Semantic Web: Research and Applications, Second European Semantic Web Conference, ESWC 2005, Heraklion, Crete, Greece, May 29 - June 1, 2005, Proceedings*, volume 3532 of *Lecture Notes in Computer Science*, pages 167–181. Springer, 2005.
49. Z. Yang, A. Ishay, and J. Lee. Neurasp: Embracing neural networks into answer set programming. In C. Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1755–1762. ijcai.org, 2020.