

Weighted Conditional \mathcal{EL}^\perp Knowledge Bases with Integer Weights: an ASP Approach

Laura Giordano

Daniele Theseider Dupré

DISIT - Università del Piemonte Orientale, Alessandria, Italy

{laura.giordano, dtd}@uniupo.it

Weighted knowledge bases for description logics with typicality have been recently considered under a “concept-wise” multipreference semantics (in both the two-valued and fuzzy case), as the basis of a logical semantics of Multilayer Perceptrons. In this paper we consider weighted conditional \mathcal{EL}^\perp knowledge bases in the two-valued case, and exploit ASP and *asprin* for encoding concept-wise multipreference entailment for weighted KBs with integer weights.

1 Introduction

Preferential approaches to common sense reasoning [26, 30, 28, 3, 23], 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 [18, 7] and closure constructions [10, 9, 19, 31, 11].

In recent work, a concept-wise multipreference semantics [16] has been proposed as a semantics of ranked \mathcal{EL}^\perp 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 rank, a natural number, representing their strength, where \mathbf{T} is a typicality operator [18] that singles out the typical instances of concept C . The concept-wise multipreference semantics takes into account preferences with respect to different concepts, and integrates them into a single global preference relation, which is needed for the evaluation of general defeasible inclusions. Answer Set Programming (ASP) and, in particular, *asprin* [6], has been exploited to achieve defeasible reasoning under the multipreference approach for the lightweight description logic \mathcal{EL}^\perp (a logic of the \mathcal{EL} family [2]).

In [17], the concept-wise multipreference semantics has been adapted to deal with weighted knowledge bases, in which each typicality inclusion is associated to a weight, a real (positive or negative) number, representing its plausibility or implausibility. The multipreference semantics has been exploited to provide a preferential interpretation of Multilayer Perceptrons (MLPs), by representing the input-output behavior of the network. It has been further extended to the fuzzy case, and the fuzzy multipreference interpretation built from the network, for a given set of input stimuli, have been proven to be a model of the neural network, in the logical sense. To this purpose, the deep network has been interpreted as a weighted conditional knowledge base, i.e., a set of weighted defeasible inclusions, by regarding synaptic connections as weighted defeasible inclusions.

While decidability of the fuzzy multipreference entailment is an open problem, in this paper we consider the two-valued case and extend the ASP approach considered for ranked KBs [16] to deal with weighted conditional KBs, for integer weights. Differently from [16], the semantic closure construction used for weighted knowledge bases does not exploit a lexicographic order, as in Lehmann lexicographic closure [28] and in Brewka’s framework of basic preference descriptions [5], but a construction more related to Kern-Isberner’s c-representations [23, 24]. We develop an ASP encoding of the concept-wise

multipreference semantics for weighted \mathcal{EL}^\perp knowledge bases, and exploit *asprin* to achieve defeasible reasoning for weighted \mathcal{EL}^\perp knowledge bases with integer weights. This is a first step towards the definition of multi-valued approximations of the fuzzy multipreference semantics for weighted knowledge bases, which may be of interest from the standpoint of explainable AI [1, 22].

2 The description logic \mathcal{EL}^\perp

We consider the description logic \mathcal{EL}^\perp of the \mathcal{EL} family [2]. 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{EL}^\perp concepts can be defined as follows: $C ::= A \mid \top \mid \perp \mid C \sqcap C \mid \exists r.C \mid \{a\}$ where $a \in N_I$, $A \in N_C$ and $r \in N_R$. $\{a\}$ is a nominal, a concept containing a single element. Observe that union, complement and universal restriction are not \mathcal{EL}^\perp constructs. 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) of the form $C \sqsubseteq D$, 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, $r \in N_R$, and $a, b \in N_I$.

An *interpretation* for \mathcal{EL}^\perp is a pair $I = \langle \Delta, \cdot^I \rangle$ where: Δ is a non-empty 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: $\top^I = \Delta$, $\perp^I = \emptyset$, $\{a\}^I = \{a^I\}$, $(C \sqcap D)^I = C^I \cap D^I$ and $(\exists r.C)^I = \{x \in \Delta \mid \exists y. (x, y) \in r^I \text{ and } y \in C^I\}$.

The notion of satisfiability of a KB in an interpretation is defined as usual:

Definition 1 (Satisfiability and Entailment) *Given an \mathcal{EL}^\perp interpretation $I = \langle \Delta, \cdot^I \rangle$:*

- *I satisfies an inclusion $C \sqsubseteq D$ if $C^I \subseteq D^I$;*

- *I satisfies an assertion $C(a)$ if $a^I \in C^I$ and an assertion $r(a, b)$ if $(a^I, b^I) \in r^I$.*

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 .

3 Weighted knowledge bases and the multipreference semantics

Let $\mathcal{C} = \{C_1, \dots, C_k\}$ be a set of distinguished \mathcal{EL}^\perp 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}^\perp 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 $C_i \in \mathcal{C}$, \mathcal{T}_{C_i} is a set of weighted 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 weighted knowledge base $K = \langle \mathcal{T}_{strict}, \mathcal{T}_{Emp}, \mathcal{T}_{Student}, \mathcal{A} \rangle$, over the set of distinguished concepts $\mathcal{C} = \{Emp, Student\}$, with empty ABox, and with \mathcal{T}_{strict} containing the set of strict inclusions:

$$Emp \sqsubseteq Adult \qquad Adult \sqsubseteq \exists has_SSN.\top \qquad PhdStudent \sqsubseteq Student$$

\mathcal{T}_{Emp} contains the following weighted defeasible inclusions:

$$(d_1) \mathbf{T}(Emp) \sqsubseteq Young, \quad -50 \qquad (d_2) \mathbf{T}(Emp) \sqsubseteq \exists has_boss.Emp, \quad 100$$

$$(d_3) \mathbf{T}(Emp) \sqsubseteq \exists has_classes.\top, \quad -70;$$

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

- (d₄) $\mathbf{T}(\text{Student}) \sqsubseteq \text{Young}, 90$ (d₅) $\mathbf{T}(\text{Student}) \sqsubseteq \exists \text{has_classes}.\top, 80$
 (d₆) $\mathbf{T}(\text{Student}) \sqsubseteq \exists \text{hasScholarship}.\top, -30$

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 (d₁) and (d₃), the second one is considered to be 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, considering the weights above, we will regard Bob as being more typical than Tom as an employee. Note that negative weights represent penalties that could not be expressed by positive ranks in a ranked \mathcal{EL}^\perp knowledge base.

The semantics of a weighted knowledge base has been defined [17] based on a concept-wise preference semantics, a semantics first exploited for ranked \mathcal{EL}^\perp knowledge bases [16]. For each concept $C_i \in \mathcal{C}$, a preference relation $<_{C_i}$ describes the preference among domain elements with respect to C_i . Each $<_{C_i}$ has the properties of preferences in KLM-style ranked interpretations [28], i.e., it is a modular and well-founded strict partial order. More precisely, $<_{C_i}$ 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 $z <_{C_j} y$).

To define a concept-wise semantics for weighted \mathcal{EL}^\perp KBs, let us recall the notion of multipreference interpretation [16]: an \mathcal{EL}^\perp interpretation is extended with a collection of preferences.

Definition 2 (Multipreference interpretation) A multipreference \mathcal{EL}^\perp 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{EL}^\perp interpretations;
 (b) for each $C_i \in \mathcal{C}$, $<_{C_i}$ is an irreflexive, transitive, well-founded and modular relation over Δ .

A 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 ($\text{tom} <_{\text{Student}} \text{bob}$), but more exceptional than Tom as an employee ($\text{bob} <_{\text{Emp}} \text{tom}$). The minimal C_i -elements with respect to $<_{C_i}$ are taken as the most typical C_i -elements.

For a multipreference interpretation \mathcal{M} to satisfy a weighted knowledge base $K = \langle \mathcal{T}_{\text{strict}}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$, we require $\langle \Delta, \cdot^I \rangle$ to satisfy the axioms in $\mathcal{T}_{\text{strict}}$ and \mathcal{A} , as usual in \mathcal{EL}^\perp , and each preference relation $<_{C_i}$ to be constructed from \mathcal{T}_{C_i} , through a *semantic closure construction*, similar in spirit to the lexicographic closure [29], but more similar to c-representation [23, 24]. The sum of the weights of the defeasible inclusions for C_i satisfied by each domain element $x \in \Delta$ is considered; higher preference wrt $<_{C_i}$ is given to the domain elements whose associated sum (wrt C_i) is higher [17].

As \mathcal{EL}^\perp has the finite model property [2], we can restrict to interpretations with a finite domain Δ . 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). Given an \mathcal{EL}^\perp interpretation $I = \langle \Delta, \cdot^I \rangle$ and a domain element $x \in \Delta$, the *weight* $W_i(x)$ of x wrt C_i in I is defined considering the inclusions $(\mathbf{T}(C_i) \sqsubseteq D_{i,h}, w_h^i) \in \mathcal{T}_{C_i}$, as follows:

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

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 defeasible inclusions for C_i satisfied by x in I . The more plausible are the satisfied inclusions, the higher is the weight of x . The lowest weight, $-\infty$, is given to all domain elements which are not instances of C_i .

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

defaults in \mathcal{T}_{C_i} satisfied by y : for $x, y \in \Delta$,

$$x <_{C_i} y \quad \text{iff} \quad W_i(x) > W_i(y) \quad (2)$$

Note that $<_{C_j}$ a strict modular partial order, and all C_i -elements are preferred wrt $<_{C_i}$ to the domain elements which are not instances of C_i . 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 $C_i = \text{Emp}$) and, hence, $\text{bob} <_{\text{Emp}} \text{tom}$, i.e., Bob is more typical than Tom as an employee.

Given the preferences $<_{C_i}$ for the distinguished concepts, we can interpret $\mathbf{T}(C_i)$ as the set of minimal C_i elements w.r.t. preference $<_{C_i}$. To provide an interpretation of the typicality concept $\mathbf{T}(C)$ for an arbitrary C (such as, for instance, $\mathbf{T}(\text{Employee} \sqcap \text{Student})$), following [16], a notion of global preference $<$ is introduced by exploiting a modified *Pareto* combination of the preference relations $<_{C_1}, \dots, <_{C_k}$, which takes into account the *specificity relation* \succ among concepts, e.g., that concept *PhDStudent* is more specific than concept *Student* ($\text{PhDStudent} \succ \text{Student}$), and its properties override the properties of *Student*, when conflicting). The global preference relation $<$ is defined from $<_{C_1}, \dots, <_{C_k}$ as follows:

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

We interpret $\mathbf{T}(C)$, for an arbitrary concept C , as the set of minimal C -elements with respect to $<$ (i.e., $(\mathbf{T}(C))^I = \min_{<}(C^I)$). This leads to the definition of a *concept-wise multipreference model* (cw^m -model) of a weighted knowledge base $K = \langle \mathcal{T}_{\text{strict}}, \mathcal{T}_{C_1}, \dots, \mathcal{T}_{C_k}, \mathcal{A} \rangle$ over \mathcal{C} , as a tuple $\mathcal{M} = \langle \Delta, <_{C_1}, \dots, <_{C_k}, <, \cdot^I \rangle$, where $\langle \Delta, \cdot^I \rangle$ satisfies $\mathcal{T}_{\text{strict}}$ and \mathcal{A} in $\mathcal{E}\mathcal{L}^\perp$; $<_{C_j}$ is defined from \mathcal{T}_{C_j} and I , according to condition (2), for all $j = 1, \dots, k$; and $<$ is the global preference relation.

Based on the notion of cw^m -model of a KB, a notion of *concept-wise entailment* (or cw^m -entailment) can be defined in a natural way for weighed KBs. Let us restrict consideration to (finite) *canonical* models, i.e., models which are large enough to contain all the relevant domain elements (see [16]).

Definition 3 (cw^m -entailment [17]) *An inclusion $\mathbf{T}(C) \sqsubseteq D$ is cw^m -entailed from a weighted knowledge base K if it is satisfied in all canonical cw^m -models \mathcal{M} of K .*

As for ranked $\mathcal{E}\mathcal{L}^\perp$ knowledge bases [16], it can be proven that this notion of cw^m -entailment for weighed $\mathcal{E}\mathcal{L}^\perp$ KBs satisfies the KLM postulates of a preferential consequence relation, as the global preference relation $<$ is a strict partial order.

4 Encoding cw^m -entailment in ASP and *asprin* for integer weights

From the computational point of view, in the case of knowledge bases integer weights, the notion of cw^m -entailment can be reformulated as a problem of computing preferred answer sets, as done for ranked knowledge bases [16]. It is possible to adapt the ASP encoding for ranked knowledge bases to deal with the extension of $\mathcal{E}\mathcal{L}^\perp$ with typicality inclusions with integer weights. The encoding has been developed by exploiting a fragment of Krötzsch's Datalog materialization calculus [27] to generate the answer sets (representing canonical cw^m -models), and *asprin* [6] to select preferred answer sets: *asprin* allows to do this according to preferences defined from a library, or with a preference program, as in this case.

In principle, verifying cw^m -entailment of a typicality subsumption $\mathbf{T}(C) \sqsubseteq D$ from K , would require considering all typical C -elements in all possible canonical cw^m -models of K , and checking whether they are all instances of D . However, it is sufficient to consider, among all the (finite) cw^m -models of K , the polynomial $\mathcal{E}\mathcal{L}^\perp$ models that we can construct using the $\mathcal{E}\mathcal{L}^\perp$ fragment of the materialization calculus by Krötzsch [27], and a distinguished domain element aux_C to represent a prototypical C -element. The

preferred answer sets are those maximizing the weight of typicality inclusions satisfied by aux_C . As in the materialization calculus auxiliary constants are used to deal with existential rules. Differently from [16], we do not need other auxiliary predicates for the distinguished concepts C_i .

Following [27], we assume that the knowledge base K is in *normal form* [2], where a typicality inclusion $\mathbf{T}(B) \sqsubseteq C$ is in normal form when $B, C \in N_C$. Extending the results by Krötzsch [27], it can be proven that, given a KB, a semantically equivalent KB in normal form (over an extended signature) can be computed in linear time [20].

The base program $\Pi(K, C, D)$ for the (normalized) knowledge base K and a typicality subsumption $\mathbf{T}(C) \sqsubseteq D$ is composed of three parts, $\Pi(K, C, D) = \Pi_K \cup \Pi_{IR} \cup \Pi_{C,D}$. Π_K is the representation of K in Datalog based on the materialization calculus [27], where $nom(a)$, $cls(A)$, $rol(R)$ are used for $a \in N_I$, $A \in N_C$, $R \in N_R$, and, for example, $subClass(a, C)$ and $subClass(A, C)$ are used, respectively, for $C(a)$ and $A \sqsubseteq C$. Additionally, $subTyp(C, D, W)$ is introduced to describe the typical properties $T(C) \sqsubseteq D$ of concept C with their weight W , and is used for specifying preferences in *asprin*.

Π_{IR} contains the subset of the inference rules for instance checking from the materialization calculus [27], those relevant for \mathcal{EL}^\perp , for example, $inst(x, z) \leftarrow subClass(y, z), inst(x, y)$. For \perp , we use an additional rule: $\leftarrow bot(z), inst(x, z)$. Additionally, Π_{IR} contains the version of the same rules for subclass checking (where $inst_sc(A, B, A)$ represents $A \sqsubseteq B$ [27]), and a rule to define predicate $morespec(Ch, Cj)$ describing the specificity relation among concepts (meaning that concept C_h is more specific than concept C_j). Π_{IR} also contains the rule: $\{inst(aux_C, D)\} \leftarrow cls(D)$, which generates alternative answer sets, corresponding to different interpretations of aux_C , which may be an instance or not of any concept D .

Let us notice that the rules enforcing \mathbf{T} -compliance in [16] are omitted. In fact, in a \mathbf{T} -compliant interpretation all typical C_i elements are assumed to satisfy *all* typicality inclusions for C_i in \mathcal{T}_{C_i} , an assumption which cannot be taken in the presence of defeasible inclusions with negative weights.

$\Pi_{C,D}$ contains (when needed) normalized axioms defining C, D in $\mathbf{T}(C) \sqsubseteq D$ in terms of other concepts (e.g., replacing $\mathbf{T}(Emp \sqcap Student) \sqsubseteq Young$ with $\mathbf{T}(A) \sqsubseteq Young$ and $A \sqsubseteq Emp$, $A \sqsubseteq Student$ and $Emp \sqcap Student \sqsubseteq A$) plus the facts $aux_C(aux_C, C)$, $nom(aux_C)$, $inst(aux_C, C)$.

Differently from the ASP encoding in [16], where the lexicographic strategy # [5] for ranked knowledge bases is used, the preference among individuals is now expressed as: $x <_{C_i} y$ iff $W_i(x) > W_i(y)$, where $W_i(x)$ and $W_i(y)$ are integer values. Given a query $\mathbf{T}(C) \sqsubseteq D$, we have to verify that, in all canonical cw^m -models of the KB, all the typical C -elements are D -elements. This verification is accomplished, by generating answer sets, corresponding to the cw^m -models of the KB, and by selecting the preferred ones, in which the distinguished element aux_C represents a typical C -element.

An answer set S is preferred to answer set S' if, considering the typicality inclusions satisfied by aux_C in S and in S' (and the relative weights $W_i(aux_C^S)$ and $W_i(aux_C^{S'})$ for each concept C_i), in the global preference relation, $aux_C^S < aux_C^{S'}$. The preference relations $<_{C_i}$ can be suitably encoded in the *asprin* specification (based on the definition above), verifying which atoms $inst(aux_C, B)$ are in S (resp., in S') to determine the typicality inclusions satisfied by aux_C^S and $aux_C^{S'}$, to establish whether $aux_C^S <_{C_i} aux_C^{S'}$. The (global) preference relation between answer sets can be defined by combining the preferences $<_{C_i}$. A query $\mathbf{T}(C) \sqsubseteq D$ is entailed by the knowledge base if, in all the preferred answer sets, aux_C is an instance of concept D . It can be proven that this corresponds to verifying that D is satisfied in all minimal C -elements in all canonical cw^m -models of the knowledge base.

Based on the reformulation of cw^m -entailment as a problem of computing preferred answer sets, deciding cw^m -entailment can be proven to be in Π_2^P for weighted knowledge bases with integer weights, as for ranked KBs. The proof of the result is similar to the proof of Proposition 7 in the online Appendix of [16], although here we use a different notion of preference $<_{C_i}$, a different encoding, and do not consider \mathbf{T} -compliant interpretations. This result also extends to weighted knowledge bases with real

valued weights, although our implementation in *asprin* does not deal with real valued weights.

5 Conclusions

In this paper we have described an ASP approach for reasoning in a defeasible extension of the description logic \mathcal{EL}^\perp with weighted typicality inclusions. We have considered the case with integer weights, for which an encoding of the concept-wise multipreference entailment can be defined using ASP and *asprin* [6]. We have further considered a concept-wise semantics with multiple preferences, in which any distinguished concept C_i has its own set \mathcal{T}_{C_i} of (weighted) typicality inclusions, and preference $<_{C_i}$.

Related semantics with multiple preferences have been proposed, starting from Brewka's framework of basic preference descriptions [5], in system ARS, as a refinement of System Z by Kern-Isberner and Ritterskamp [25]; in an extension of $\mathcal{ALC} + \mathbf{T}$ by Gil [14]; in a refinement of rational closure by Gliozzi [21]; by associating multiple preferences to roles by Britz and Varzinczak [8]; in ranked \mathcal{EL} knowledge bases in [16]; and in the first-order logic setting by Delgrande and Rantsaudis [13].

The development of proof methods for multipreference entailment is motivated by the relationship between weighted conditionals under the (fuzzy) concept-wise multipreference semantics and multilayer perceptrons. Undecidability results for fuzzy description logics with general inclusion axioms [12, 4] motivate the investigation of multi-valued approximations of fuzzy multipreference entailment. This is a first step towards the definition of proof methods for multi-valued extensions of our concept-wise preferential semantics based on a notion of faithful interpretations [15]. Other possible extensions concern the definition of multiple typicality operators, based on the combination of selected concepts, and a temporal extension to capture the transient behavior of Multilayer Perceptrons.

Acknowledgement: We thank the referees for their helpful suggestions. This research is partially supported by INDAM-GNCS Project 2020.

References

- [1] A. Adadi & M. Berrada (2018): *Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)*. *IEEE Access* 6, pp. 52138–52160, doi:10.1109/ACCESS.2018.2870052.
- [2] F. Baader, S. Brandt & C. Lutz (2005): *Pushing the \mathcal{EL} envelope*. In: *Proc. 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, Edinburgh, Scotland, UK, pp. 364–369.
- [3] S. Benferhat, C. Cayrol, D. Dubois, J. Lang & H. Prade: *Inconsistency Management and Prioritized Syntax-Based Entailment*. In: *Proc. IJCAI'93, Chambéry, France, 28/8 - 3/9, 1993*, Morgan Kaufmann, pp. 640–647.
- [4] S. Borgwardt & R. Peñaloza (2012): *Undecidability of Fuzzy Description Logics*. In: *Principles of Knowledge Representation and Reasoning: Proc. 13th Int. Conf., KR 2012, Rome, Italy, June 10-14, 2012*.
- [5] G. Brewka (2004): *A Rank Based Description Language for Qualitative Preferences*. In: *Proc. 16th European Conf. on Artificial Intelligence, ECAI'2004, Valencia, Spain, August 22-27, 2004*, pp. 303–307.
- [6] G. Brewka, J. P. Delgrande, J. Romero & T. Schaub (2015): *asprin: Customizing Answer Set Preferences without a Headache*. In: *Proc. AAI 2015*, pp. 1467–1474.
- [7] K. Britz, J. Heidema & T. Meyer (2008): *Semantic Preferential Subsumption*. In G. Brewka & J. Lang, editors: *KR 2008*, AAI Press, Sidney, Australia, pp. 476–484.
- [8] K. Britz & I. J. Varzinczak (2018): *Rationality and Context in Defeasible Subsumption*. In: *FoIKS 2018, Budapest, May 14-18, 2018*, pp. 114–132, doi:10.1007/978-3-319-90050-6_7.
- [9] G. Casini, T. Meyer, I. J. Varzinczak, & K. Moodley (2013): *Nonmonotonic Reasoning in Description Logics: Rational Closure for the ABox*. In: *DL 2013, CEUR Workshop Proceedings 1014*, pp. 600–615.

- [10] G. Casini & U. Straccia (2010): *Rational Closure for Defeasible Description Logics*. In: *JELIA 2010, LNCS 6341*, Springer, Helsinki, pp. 77–90, doi:10.1007/978-3-642-15675-5_9.
- [11] G. Casini, U. Straccia & T. Meyer (2019): *A polynomial Time Subsumption Algorithm for Nominal Safe $\mathcal{E}\mathcal{L}\perp$ under Rational Closure*. *Inf. Sci.* 501, pp. 588–620, doi:10.1016/j.ins.2018.09.037.
- [12] M. Cerami & U. Straccia (2011): *On the Undecidability of Fuzzy Description Logics with GCIs with Lukasiewicz t-norm*. CoRR abs/1107.4212. Available at <http://arxiv.org/abs/1107.4212>.
- [13] J. Delgrande & C. Rantsoudis (2020): *A Preference-Based Approach for Representing Defaults in First-Order Logic*. In: *18th Int. Workshop on Non-Monotonic Reasoning, NMR2020, September 12th - 14th*.
- [14] Oliver Fernandez Gil (2014): *On the Non-Monotonic Description Logic $ALC+T_{min}$* . CoRR abs/1404.6566.
- [15] L. Giordano (2021): *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, Sept. 17-20, LNAI 12897*, Springer, doi:10.1007/978-3-030-86772-0_40.
- [16] L. Giordano & D. Theseider Dupré (2020): *An ASP approach for reasoning in a concept-aware multipreferential lightweight DL*. *TPLP* 10(5), pp. 751–766, doi:10.1007/978-3-642-15675-5_21.
- [17] L. Giordano & D. Theseider Dupré (2021): *Weighted defeasible knowledge bases and a multipreference semantics for a deep neural network model*. In: *Proc 17th European Conf. on Logics in AI, JELIA 2021, May 17-20, LNCS 12678*, Springer, pp. 225–242, doi:10.1007/978-3-030-75775-5_16.
- [18] L. Giordano, V. Gliozzi, N. Olivetti & G. L. Pozzato (2007): *Preferential Description Logics*. In: *LPAR 2007, LNAI 4790*, Springer, Yerevan, Armenia, pp. 257–272, doi:10.1007/978-3-540-75560-9_20.
- [19] L. Giordano, V. Gliozzi, N. Olivetti & G. L. Pozzato (2015): *Semantic characterization of rational closure: From propositional logic to description logics*. *Artif. Intell.* 226, pp. 1–33, doi:10.1016/j.artint.2015.05.001.
- [20] L. Giordano & D. Theseider Dupré (2018): *Defeasible Reasoning in SROEL from Rational Entailment to Rational Closure*. *Fundam. Inform.* 161(1-2), pp. 135–161, doi:10.3233/FI-2018-1698.
- [21] V. Gliozzi (2016): *Reasoning about Multiple Aspects in Rational Closure for DLs*. In: *AI*IA 2016: 15th Int. Conf. of the Italian Assoc. for AI, Genova, Italy, Nov. 2016*, pp. 392–405, doi:10.1007/978-3-319-49130-1_29.
- [22] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti & D. Pedreschi (2019): *A Survey of Methods for Explaining Black Box Models*. *ACM Comput. Surv.* 51(5), pp. 93:1–93:42, doi:10.1145/3236009.
- [23] G. Kern-Isberner (2001): *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*. LNCS 2087, Springer, doi:10.1007/3-540-44600-1.
- [24] G. Kern-Isberner & C. Eichhorn (2014): *Structural Inference from Conditional Knowledge Bases*. *Studia Logica* 102(4), pp. 751–769, doi:10.1007/s11225-013-9503-6.
- [25] G. Kern-Isberner & M. Ritterskamp (2010): *Preference Fusion for Default Reasoning Beyond System Z*. *J. Autom. Reasoning* 45(1), pp. 3–19, doi:10.1007/s10817-009-9129-6.
- [26] S. Kraus, D. Lehmann & M. Magidor (1990): *Nonmonotonic Reasoning, Preferential Models and Cumulative Logics*. *Artificial Intelligence* 44(1-2), pp. 167–207, doi:10.1016/0004-3702(90)90101-5.
- [27] M. Krötzsch (2010): *Efficient Inferencing for OWL EL*. In: *Proc. JELIA 2010*, pp. 234–246, doi:10.1007/978-3-642-15675-5_21.
- [28] D. Lehmann & M. Magidor (1992): *What does a conditional knowledge base entail?* *Artificial Intelligence* 55(1), pp. 1–60, doi:10.1016/0004-3702(92)90041-U.
- [29] D. J. Lehmann (1995): *Another Perspective on Default Reasoning*. *Ann. Math. Artif. Intell.* 15(1), pp. 61–82, doi:10.1007/BF01535841.
- [30] J. Pearl: *System Z: A Natural Ordering of Defaults with Tractable Applications to Nonmonotonic Reasoning*. In: *TARK'90, Pacific Grove, CA, USA, 1990*, Morgan Kaufmann, pp. 121–135.
- [31] M. Pensel & A. Turhan (2018): *Reasoning in the Defeasible Description Logic EL_\perp - computing standard inferences under rational and relevant semantics*. *Int. J. Approx. Reasoning* 103, pp. 28–70, doi:10.1016/j.ijar.2018.08.005.