

FUZZY OWL-BOOST: Learning Fuzzy Concept Inclusions via Real-Valued Boosting

Franco Alberto Cardillo
ILC - CNR
Pisa, Italy

Umberto Straccia
ISTI - CNR
Pisa, Italy

March 29, 2021

Abstract

OWL ontologies are nowadays a quite popular way to describe structured knowledge in terms of classes, relations among classes and class instances.

In this paper, given an OWL ontology and a target class T , we address the problem of learning fuzzy concept inclusion axioms that describe sufficient conditions for being an individual instance of T (and to which degree). To do so, we present FUZZY OWL-BOOST that relies on the \mathbb{R} Real AdaBoost boosting algorithm adapted to the (fuzzy) OWL case. We illustrate its effectiveness by means of an experimentation with several ontologies.

1 Introduction

OWL 2 ontologies [66] are nowadays a popular means to represent *structured* knowledge and its formal semantics is based on *Description Logics* (DLs) [4]. The basic ingredients of DLs are concept descriptions (in First-Order Logic terminology, unary predicates), inheritance relationships among them and instances of them.

Although an important amount of work has been carried about DLs, the application of machine learning techniques to OWL 2 ontologies, *viz.* DL ontologies, is relatively less addressed compared to the *Inductive Logic Programming* (ILP) setting (see *e.g.* [68, 69] for more insights on ILP). We refer the reader to [51, 70] for an overview.

In this work, we focus on the problem of automatically learning fuzzy concept inclusion axioms from (crisp) OWL 2 ontologies. More specifically, given a target class T of an OWL ontology, we address the problem of learning fuzzy concept inclusion axioms that describe sufficient conditions for being an individual instance of T (and to which degree). An example illustrating the problem is shown next.

Example 1.1 (Running example [48, 50, 85]) *Consider an ontology that describes the meaningful entities of a city, such as e.g. the Hotel ontology in our experiments (see Section 5, Table 2). An excerpt of this ontology is given in Fig. 1.*

Now, one may fix a city, say Pisa, extract the properties of the hotels from Web sites, such as location, price, etc., and the hotel judgements of the users, e.g., from Trip Advisor.¹ Now, using the terminology of the ontology and class instances gathered from the Web one may ask about

¹<http://www.tripadvisor.com>

Let us remark that we rely on real-valued AdaBoost as the weak hypotheses FUZZY OWL-BOOST generates are indeed fuzzy concept inclusion axioms and, thus, the degree to which an instance satisfies them is a real-valued degree of truth in $[0, 1]$.

In the following, we proceed as follows. In Section 2 we compare our work with closely related work appeared so far. In Section 3, for the sake of completeness, we recap the salient notions we will rely on in this paper. Then, in Section 4 we will present our algorithm FUZZY OWL-BOOST that then is evaluated for its effectiveness in Section 5. Section 6 concludes and points to some topics of further research.

2 Related Work

Concept inclusion axiom learning in DLs stems from statistical relational learning, where classification rules are (possibly weighted) Horn clause theories (see *e.g.* [68, 69]), and various methods have been proposed in the DL context so far (see *e.g.* [51, 70]). The general idea consists in the exploration of the search space of potential concept descriptions that cover the available training examples using so-called refinement operators (see, *e.g.* [5, 19, 42, 43, 44, 45, 46]). The goal is then to learn a concept description of the underlying DL language covering (possibly) all the provided positive examples and (possibly) not covering any of the provided negative examples. The fuzzy case (see [47, 50, 85]) is a natural extension relying on fuzzy DLs [9, 84] and fuzzy ILP (see *e.g.* [80]) instead.

Closely related to our work are [26, 24, 32, 31, 47, 50, 85]. In fact, [26, 24, 32] are an adaptation to the DL case of the the well-known FOIL-algorithm, while [47, 50] that stem essentially from [48, 49, 52, 53, 54, 55], propose *fuzzy* FOIL-like algorithms instead, and are inspired by fuzzy ILP variants such as [22, 80, 82].³ Let us note that [47, 53] consider the weaker hypothesis representation language DL-Lite [2], while here we rely on a weighted sum of fuzzy $\mathcal{EL}(\mathbf{D})$ inclusion axioms, similarly to [48, 49, 52, 54, 55, 50]. Fuzzy $\mathcal{EL}(\mathbf{D})$ has also been considered in [85], which however differs from [47, 50] by the fact that a (fuzzy) probabilistic ensemble evaluation of the fuzzy concept description candidates has been considered.⁴ Let us note that fuzzy $\mathcal{EL}(\mathbf{D})$ concept expressions are appealing as they can straightforwardly be translated into natural language and, thus, contribute to the explainability aspect of the induced classifier.

Discrete boosting has been considered in [31] that also shows how to derive a weak learner (called WDLF) from conventional learners using some sort of random downward refinement operator covering at least a positive example and yielding a minimal score fixed with a threshold. Besides that, we deal here with fuzziness in the hypothesis language and a real-valued variant of AdaBoost, the weak learner we propose here differentiates from the previous one by using a descent-like gradient algorithm to search for the best alternative. Notably, this also deviates from ‘fuzzy’ rule learning AdaBoost variants, such as [21, 65, 67, 78, 87] in which the weak learner is required to generate the whole rules’ search space beforehand the selection of the best current alternative. Such an approach is essentially unfeasible in the OWL case due to the size of the search space.

[37] can learn fuzzy OWL DL concept equivalence axioms from FuzzyOWL 2 ontologies, by interfacing with the *fuzzyDL* reasoner [13]. The candidate concept expressions are provided by the underlying DL-LEARNER [41, 16, 17] system. However, it has been tested only on a toy ontology so far. Last, but not least, let us mention [39] that is based on an ad-hoc translation of fuzzy Łukasiewicz \mathcal{ALC} DL constructs into fuzzy *Logic Programming* (fuzzy LP) and uses a

concepts we are going to use. The extension is straightforward.

³See, *e.g.* [20], for an overview on fuzzy rule learning methods.

⁴Also, to the best of our knowledge, concrete datatypes were not addressed in the evaluation.




Figure 2: (a) Trapezoidal function $trz(a, b, c, d)$, (b) triangular function $tri(a, b, c)$, (c) left shoulder function $ls(a, b)$, and (d) right shoulder function $rs(a, b)$.

conventional ILP method to learn rules. Unfortunately, the method is not sound as it has been shown that the mapping from fuzzy DLs to LP is incomplete [61] and entailment in Łukasiewicz \mathcal{ALC} is undecidable [18]. To be more precise, undecidability holds already for \mathcal{EL} under the infinitely valued Łukasiewicz semantics [15].⁵

While it is not our aim here to provide an extensive overview about learning w.r.t. ontologies literature, there are also alternative methods to what we present here. So, *e.g.*, the series of works [28, 29, 73, 72, 74, 71, 77, 75, 76] are inspired on *Decision Trees/Random Forests*, [8, 25, 27, 30] consider *Kernel Methods* for inducing concept descriptions, while [57, 59, 58, 60, 89] consider essentially a *Naive Bayes* approach. Last but not least, [40] is inspired on *Genetic Programming* to induce concept expressions, while [62] is based on the *Reinforcement Learning* framework.

3 Background

For the sake of completeness, we recap here the salient notions about *fuzzy Description Logics* (fuzzy DLs) we will rely on in this paper. The interested reader may refer to *e.g.* [9, 84] for a more in depth description of the various notions introduced here.

Fuzzy Sets To start with, we recap that a *fuzzy set* A over a countable crisp set X is a function $A: X \rightarrow [0, 1]$, called *fuzzy membership function* of A [88]. A crisp set A is defined by a membership function $A: X \rightarrow \{0, 1\}$ instead. The ‘standard’ fuzzy set operations conform to $(A \cap B)(x) = \min(A(x), B(x))$, $(A \cup B)(x) = \max(A(x), B(x))$ and $\bar{A}(x) = 1 - A(x)$ (\bar{A} is the set complement of A), the *cardinality* of a fuzzy set is defined as

$$|A| = \sum_{x \in X} A(x), \quad (1)$$

while the *inclusion degree* of A in B is defined as

$$ideg(A, B) = \frac{|A \cap B|}{|A|}. \quad (2)$$

⁵We recall that \mathcal{EL} is a strict sub-logic of \mathcal{ALC} .




Figure 3: Uniform fuzzy sets over salaries.

The trapezoidal (Fig. 2 (a)), the triangular (Fig. 2 (b)), the L -function (left-shoulder function, Fig. 2 (c)), and the R -function (right-shoulder function, Fig. 2 (d)) are frequently used to specify membership functions of fuzzy sets.

Although fuzzy sets have a greater expressive power than classical crisp sets, their usefulness depends critically on the capability to construct appropriate membership functions for various given concepts in different contexts. We refer the interested reader to, *e.g.*, [38]. One easy and typically satisfactory method to define the membership functions is to uniformly partition the range of, *e.g.* salary values (bounded by a minimum and maximum value), into 3, 5 or 7 fuzzy sets using triangular (or trapezoidal) functions (see Fig. 3). Another popular approach may consist in using the so-called C -means fuzzy clustering algorithm (see, *e.g.* [7]) with 3, 5 or 7 clusters, where the fuzzy membership functions are triangular functions built around the centroids of the clusters (see *e.g.* also [35, 36]).

Fuzzy Description Logics Next, we recap here the fuzzy DL $\mathcal{ELW}^\top(\mathbf{D})$ extending the well-known fuzzy DL \mathcal{EL} [9, 84] with⁶

- *real-valued weighted concept* construct (denoted by the letter \mathcal{W}) [12, 84];
- *atomic negation* (denoted by \neg);
- *fuzzy concrete domains* (denoted by \mathbf{D}) [83].

Fuzzy $\mathcal{ELW}^\top(\mathbf{D})$ is expressive enough to capture all of the ingredients we are going to use in this work. Of course, DLs, fuzzy DLs, OWL 2 and fuzzy OWL 2 in particular, cover many more language constructs than we use here (see, *e.g.* [4, 9, 12, 84]).

From a syntax point of view, we start with the notion of *fuzzy concrete domain*, that is a tuple $\mathbf{D} = \langle \Delta^{\mathbf{D}}, \cdot^{\mathbf{D}} \rangle$ with datatype domain $\Delta^{\mathbf{D}}$ and a mapping $\cdot^{\mathbf{D}}$ that assigns to each data value an element of $\Delta^{\mathbf{D}}$, and to every 1-ary datatype predicate \mathbf{d} (see below) a 1-ary fuzzy relation over $\Delta^{\mathbf{D}}$. Therefore, $\cdot^{\mathbf{D}}$ maps indeed each datatype predicate \mathbf{d} into a function from $\Delta^{\mathbf{D}}$ to $[0, 1]$. The domains we consider here are the integers, the reals and the booleans. The datatype predicates are defined as

$$\mathbf{d} \rightarrow ls(a, b) \mid rs(a, b) \mid tri(a, b, c) \mid trz(a, b, c, d) \mid =_v ,$$

⁶For classical \mathcal{EL} and DLs in general we refer the reader to [3, 4].

where v is an integer, real or boolean value, *e.g.* $ls(a, b)$ is the left-shoulder membership function (see Fig. 2) and $=_v$ corresponds to the crisp singleton set $\{v\}$.

Now, consider pairwise disjoint alphabets \mathbf{I} , \mathbf{A} and \mathbf{R} , where \mathbf{I} is the set of *individuals*, \mathbf{A} is the set of *concept names* (also called *atomic concepts*) and \mathbf{R} is the set of *role names*. Each role is either an *object property* or a *datatype property*. The set of fuzzy $\mathcal{ELW}^\neg(\mathbf{D})$ *concepts* are built from concept names A using connectives and quantification constructs over object properties r and datatype properties s , as described by the following syntactic rule ($n \geq 1, \alpha_i \in \mathbb{R} \setminus \{0\}$):

$$C \rightarrow \top \mid \perp \mid A \mid \neg A \mid C_1 \sqcap C_2 \mid \exists r.C \mid \mid \exists s.d \mid \alpha_1 \cdot A_1 + \dots + \alpha_n \cdot A_n .$$

Note that we generalise slightly the notion of weighted sum of Fuzzy OWL 2 in which $\alpha_i \in (0, 1], \sum_i \alpha_i \leq 1$ is assumed.

A *fuzzy assertion* axiom is an expression of the form $\langle a:C, d \rangle$ (called *fuzzy concept assertion*, $-a$ is an instance of concept C to degree greater than or equal to d) or of the form $\langle (a_1, a_2):r, d \rangle$ (called *fuzzy role assertion*, $-(a_1, a_2)$ is an instance of object property r to degree greater than or equal to d), where a, a_1, a_2 are individual names, C is a concept, r is an object property and $d \in (0, 1]$.

A *fuzzy General Concept Inclusion* (fuzzy GCI) axiom is of the form $\langle C_1 \sqsubseteq C_2, d \rangle$ (C_1 is a sub-concept of C_2 to degree greater than or equal to d), where C_i is a concept and $d \in (0, 1]$. We may also call a fuzzy GCI of the form $\langle C \sqsubseteq A, d \rangle$, where A is a concept name, a *rule*, where A is called the *head* and C is called the *body* of the rule.

For ease of presentation, in case the degree d is 1, we simply omit the degree 1 and write, *e.g.* $a:C$ (resp. $C_1 \sqsubseteq C_2$) in place of $\langle a:C, 1 \rangle$ (resp. $\langle C_1 \sqsubseteq C_2, 1 \rangle$). We also write $C_1 = C_2$ as a macro for the two GCIs $C_1 \sqsubseteq C_2$ and $C_2 \sqsubseteq C_1$, indicating that the two concepts C_1 and C_2 are equivalent.

Fuzzy $\mathcal{EL}(\mathbf{D})$ is the language fuzzy $\mathcal{ELW}^\neg(\mathbf{D})$ without atomic negation and weighted sum, while *crisp $\mathcal{EL}(\mathbf{D})$* , or simply $\mathcal{EL}(\mathbf{D})$, is fuzzy $\mathcal{EL}(\mathbf{D})$ without *ls, rs, tri* and *trz* datatype predicates and the degree in axioms is restricted to be 1.

A (crisp) *Knowledge Base* (KB) \mathcal{K} is a set of (crisp) $\mathcal{EL}(\mathbf{D})$ assertions and GCIs, while a *fuzzy Knowledge Base* (fKB) $\tilde{\mathcal{K}}$ is a set of fuzzy $\mathcal{ELW}^\neg(\mathbf{D})$ assertions and GCIs.

Remark 1 *We anticipate that in our setting we are going to learn fuzzy GCIs from crisp OWL 2 data, i.e. we are going to learn from crisp knowledge bases only.*

With $l_{\mathcal{K}}$ we denote the set of individuals occurring in a KB \mathcal{K} .

Example 3.1 (Example 1.1 cont.) *Related to Example 1.1, an example of GCI is*

$$\text{Hotel} \sqsubseteq \text{Accommodation}$$

(a hotel is an accommodation), while the assertions

$$\begin{aligned} & (\text{Hotel}_010, 3_Stars):hasRank \\ & \text{Hotel}_010:\text{Hotel} \sqcap (\exists hasAmenity.24h_Reception) \sqcap (\exists hasPrice. =_{79}) \end{aligned}$$

describe a three star hotel (Hotel_010) that provides 24h reception as amenity and whose room price is 79 (euro). \square

Concerning the semantics, let us fix a fuzzy concrete domain $\mathbf{D} = \langle \Delta^{\mathbf{D}}, \cdot^{\mathbf{D}} \rangle$ over integer, reals and booleans, with fuzzy membership functions *ls, rs, tri* and *trz* over integers and reals, and equality predicate $=$ over integers, reals and booleans.

Now, unlike classical DLs in which an interpretation \mathcal{I} maps *e.g.* a concept C into a set of individuals $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, *i.e.* \mathcal{I} maps C into a function $C^{\mathcal{I}} : \Delta^{\mathcal{I}} \rightarrow \{0, 1\}$ (either an individual belongs to the extension of C or does not belong to it), in fuzzy DLs, \mathcal{I} maps C into a function $C^{\mathcal{I}} : \Delta^{\mathcal{I}} \rightarrow [0, 1]$ and, thus, an individual belongs to the extension of C to some degree in $[0, 1]$, *i.e.* $C^{\mathcal{I}}$ is a fuzzy set.

Specifically, a *fuzzy interpretation* is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consisting of a nonempty (crisp) set $\Delta^{\mathcal{I}}$ (the *domain*) and of a *fuzzy interpretation function* $\cdot^{\mathcal{I}}$ that assigns: (i) to each atomic concept A a function $A^{\mathcal{I}} : \Delta^{\mathcal{I}} \rightarrow [0, 1]$; (ii) to each object property r a function $r^{\mathcal{I}} : \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \rightarrow [0, 1]$; (iii) to each datatype property s a function $s^{\mathcal{I}} : \Delta^{\mathcal{I}} \times \Delta^{\mathbf{D}} \rightarrow [0, 1]$; (iv) to each individual a an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ such that $a^{\mathcal{I}} \neq b^{\mathcal{I}}$ if $a \neq b$ (the so-called *Unique Name Assumption*); and (v) to each data value v an element $v^{\mathcal{I}} \in \Delta^{\mathbf{D}}$. Now, a fuzzy interpretation function is extended to concepts via standard fuzzy set operations as specified below (where $x \in \Delta^{\mathcal{I}}$):⁷

$$\begin{aligned}
\top^{\mathcal{I}}(x) &= 1 \\
\perp^{\mathcal{I}}(x) &= 0 \\
(C \sqcap D)^{\mathcal{I}}(x) &= \min(C^{\mathcal{I}}(x), D^{\mathcal{I}}(x)) \\
(\neg A)^{\mathcal{I}}(x) &= 1 - A^{\mathcal{I}}(x) \\
(\exists r.C)^{\mathcal{I}}(x) &= \sup_{y \in \Delta^{\mathcal{I}}} \{\min(r^{\mathcal{I}}(x, y), C^{\mathcal{I}}(y))\} \\
(\exists s.\mathbf{d})^{\mathcal{I}}(x) &= \sup_{y \in \Delta^{\mathbf{D}}} \{\min(s^{\mathcal{I}}(x, y), \mathbf{d}^{\mathbf{D}}(y))\} \\
(\alpha_1 \cdot A_1 + \dots + \alpha_n \cdot A_n)^{\mathcal{I}}(x) &= \begin{cases} 1 & \text{if } \sum_i \alpha_i \cdot A_i^{\mathcal{I}}(x) > 1 \\ 0 & \text{if } \sum_i \alpha_i \cdot A_i^{\mathcal{I}}(x) < 0 \\ \sum_i \alpha_i \cdot A_i^{\mathcal{I}}(x) & \text{else .} \end{cases}
\end{aligned}$$

The *satisfiability of axioms* is defined by the following conditions: (i) \mathcal{I} satisfies $\langle a:C, d \rangle$ if $C^{\mathcal{I}}(a^{\mathcal{I}}) \geq d$; (ii) \mathcal{I} satisfies $\langle (a, b):r, d \rangle$ if $r^{\mathcal{I}}(a^{\mathcal{I}}, b^{\mathcal{I}}) \geq d$; and (iii) \mathcal{I} satisfies $\langle C \sqsubseteq D, d \rangle$ if for all $x \in \Delta^{\mathcal{I}}$, $D^{\mathcal{I}}(x) \geq C^{\mathcal{I}}(x) \cdot d$.

Now, consider a set \mathcal{S} of axioms. Then, (i) \mathcal{I} is a *model* of \mathcal{S} if \mathcal{I} satisfies each axiom in \mathcal{S} ; (ii) \mathcal{S} is *satisfiable* (or *consistent*) if \mathcal{S} has a model; (iii) \mathcal{S} *entails* axiom ϕ , denoted $\mathcal{S} \models \phi$, if every model of \mathcal{S} satisfies ϕ ; and (iv) the *best entailment degree* of ϕ of the form $C \sqsubseteq D$, $a:C$ or $(a, b):r$, denoted $bed(\mathcal{S}, \phi)$, is defined as

$$bed(\mathcal{S}, \phi) = \sup\{d \mid \mathcal{S} \models \langle \phi, d \rangle\} .$$

Example 3.2 (Example 3.1 cont.) Consider a KB \mathcal{K} whose excerpt is described in Example 3.1. Consider the following fuzzy GCI ϕ

$$\langle \text{Accommodation} \sqcap (\exists \text{hasPrice.Fair}) \sqsubseteq \text{GoodHotel}, 0.56 \rangle ,$$

where *hasPrice* is a datatype property whose values are measured in euros and the price concrete domain has been fuzzified as illustrated in Fig. 4. The intended meaning of this axiom is roughly ‘an accommodation with a fair room price is to some degree a good hotel’. Also consider *Hotel_010* in Example 3.1.

Now, it can be verified that

$$\mathcal{K} \cup \{\phi\} \models \langle \text{Hotel_010:GoodHotel}, 0.28 \rangle ,$$

⁷The semantics of the weighted sum will be clearer once we address the learning problem.




Figure 4: Fuzzy sets derived from the datatype property `hasPrice`.

and, more specifically that

$$bed(\mathcal{K} \cup \{\phi\}, \text{Hotel_010:GoodHotel}) = 0.28 ,$$

which indicates to which degree `Hotel_010` is a `GoodHotel`.

□

Finally, consider a fuzzy concept C , a fuzzy GCI $C \sqsubseteq D$, a KB \mathcal{K} , a set of individuals I and a (weight) distribution \mathbf{w} over I . Then the *cardinality* of C w.r.t. \mathcal{K} and I , denoted $|C|_{\mathcal{K}}^I$, is defined as (cf. Eq. 1)

$$|C|_{\mathcal{K}}^I = \sum_{a \in I} bed(\mathcal{K}, a:C) , \quad (3)$$

while the *weighted cardinality* C w.r.t. \mathcal{K} , \mathbf{w} and I , denoted $|C|_{\mathcal{K}}^{\mathbf{w},I}$, is defined as⁸

$$|C|_{\mathcal{K}}^{\mathbf{w},I} = \sum_{a \in I} w_a \cdot bed(\mathcal{K}, a:C) . \quad (4)$$

The *crisp cardinality* (denoted $\lceil C \rceil_{\mathcal{K}}^I$) and *crisp weighted cardinality* (denoted $\lceil C \rceil_{\mathcal{K}}^{\mathbf{w},I}$) are defined similarly by replacing in Eq. 3 and 4 the term $bed(\mathcal{K}, a:C)$ with $\lceil bed(\mathcal{K}, a:C) \rceil$.

Furthermore, the *confidence degree* (also called *inclusion degree*) of $C \sqsubseteq D$ w.r.t. \mathcal{K} and I , denoted $cf(C \sqsubseteq D, I)$, is defined as (cf. Eq. 2)

$$cf(C \sqsubseteq D, I) = \frac{|C \sqcap D|_{\mathcal{K}}^I}{|C|_{\mathcal{K}}^I} . \quad (5)$$

Similarly, the *weighted confidence degree* (also called *weighted inclusion degree*) of $C \sqsubseteq D$ w.r.t. \mathcal{K} , \mathbf{w} and I , denoted $cf(C \sqsubseteq D, \mathbf{w}, I)$, is defined as

$$cf(C \sqsubseteq D, \mathbf{w}, I) = \frac{|C \sqcap D|_{\mathcal{K}}^{\mathbf{w},I}}{|C|_{\mathcal{K}}^{\mathbf{w},I}} . \quad (6)$$

4 Learning Fuzzy Concept Inclusions via Real-Valued Boosting

To start with, we introduce our learning problem.

⁸The weight of $a \in I$ w.r.t. \mathbf{w} is denoted as w_a .

4.1 The Learning Problem

In general terms, the learning problem we are going to address is stated as follows:

Given:

- a satisfiable (crisp) KB \mathcal{K} and its individuals $\mathsf{I}_{\mathcal{K}}$;
- a *target concept name* T with an associated unknown classification function $f_T: \mathsf{I}_{\mathcal{K}} \rightarrow \{1, 0\}$, where for each $a \in \mathsf{I}_{\mathcal{K}}$, the possible values (*labels*) correspond, respectively, to $\mathcal{K} \models a:T$ (a is a *positive* example of T) and $\mathcal{K} \not\models a:T$ (a is a *non-positive* example of T);
- a *hypothesis space* of classifiers $\mathcal{H} = \{h: \mathsf{I}_{\mathcal{K}} \rightarrow [0, 1]\}$;
- a *training set* $\mathcal{E} = \mathcal{E}^+ \cup \mathcal{E}^-$ (the positive and non-positive examples of T , respectively) of individual-label pairs:

$$\begin{aligned} \mathcal{E}^+ &= \{(a, 1) \mid a \in \mathsf{I}_{\mathcal{K}}, f_T(a) = 1\} \\ \mathcal{E}^- &= \{(a, 0) \mid a \in \mathsf{I}_{\mathcal{K}}, f_T(a) = 0\} . \end{aligned}$$

With $\mathsf{I}_{\mathcal{E}}$ we denote the set of individuals occurring in \mathcal{E} . We assume that for all $a \in \mathsf{I}_{\mathcal{E}}$, $0 = \text{bed}(\mathcal{K}, a:T) = \text{bed}(\mathcal{K}, a:\neg T)$. That is we state that \mathcal{K} does not already know whether a is an instance of T or of $\neg T$. We write $\mathcal{E}(a) = 1$ if a is a positive example (*i.e.*, $a \in \mathsf{I}_{\mathcal{E}^+}$), $\mathcal{E}(a) = 0$ if a is a non-positive example (*i.e.*, $a \in \mathsf{I}_{\mathcal{E}^-}$).

Learn: a classifier $\bar{h} \in \mathcal{H}$ that is the result of *Empirical Risk Minimisation* (ERM) on \mathcal{E} . That is,

$$\begin{aligned} \bar{h} &= \arg \min_{h \in \mathcal{H}} R(h, \mathcal{E}) \\ &= \mathbf{E}_{\mathcal{E}}[L(h(a), \mathcal{E}(a))] \\ &= \frac{1}{|\mathcal{E}|} \sum_{a \in \mathsf{I}_{\mathcal{E}}} L(h(a), \mathcal{E}(a)) , \end{aligned}$$

where L is a *loss function* such that $L(\hat{l}, l)$ measures how different the prediction \hat{l} of a hypothesis is from the true outcome l and $R(h, \mathcal{E})$ is the *risk* associated with hypothesis h over \mathcal{E} , defined as the expectation of the loss function over \mathcal{E} .

The effectiveness of the learned classifier \bar{h} is then assessed by determining $R(\bar{h}, \mathcal{E}')$ on a *test set* \mathcal{E}' , disjoint from \mathcal{E} .

In our learning setting, we assume that a hypothesis $h \in \mathcal{H}$ is a set of fuzzy GCIs that has the form

$$\begin{aligned} \alpha_1 \cdot WL_1 + \dots + \alpha_n \cdot WL_n &\sqsubseteq T \\ C_{ij} &\sqsubseteq WL_i , \text{ with } 1 \leq i \leq n, 1 \leq j \leq k_i , \end{aligned} \quad (7)$$

where each WL_i is a new atomic symbol not occurring in the KB and where each C_{ij} is a fuzzy $\mathcal{EL}(\mathbf{D}^-)$ concept expression defined as (where v is a boolean value)

$$\begin{aligned} C &\rightarrow \top \mid A \mid \exists r.C \mid \exists s.\mathbf{d} \mid C_1 \sqcap C_2 \\ \mathbf{d} &\rightarrow ls(a, b) \mid rs(a, b) \mid tri(a, b, c) \mid trz(a, b, c, d) \mid =_v . \end{aligned}$$

Essentially, each α_i indicates how well the ‘union’ of the C_{i1}, \dots, C_{ik_i} (the *Weak Learner* WL_i) contributes to classify an individual a as being an instance of T . Specifically, if $\alpha_i > 0$ then WL_i contributes to a ’s positiveness, while if $\alpha_i < 0$ then WL_i contributes to a ’s non-positiveness instead.

Remark 2 Please note that we do not learn expressions of the form e.g. $\exists s. =_v$ for integer/real values v as the search space would be too large and they would be likely non-effective. This is the reason why we restrict the C_{ij} in Eq. 7 to fuzzy $\mathcal{EL}(\mathbf{D}^-)$ concept expressions and not fuzzy $\mathcal{EL}(\mathbf{D})$ instead.

For $a \in \mathcal{I}_{\mathcal{K}}$, the classification prediction value $h(a)$ of a w.r.t. h , T and \mathcal{K} is defined as (for ease, we omit \mathcal{K} and T)

$$h(a) = \text{bed}(\mathcal{K} \cup h, a:T) .$$

Note that, as stated above, essentially a hypothesis is a sufficient condition (expressed via the weighted sum of concepts) for being an individual instance of a target concept to some degree. If $h(a) = 0$ then we say that a is a non-positive instance of T , while if $h(a) > 0$ then a is a positive instance of T to degree $h(a)$.

Remark 3 Clearly, the set of hypothesis by this syntax is potentially infinite due, e.g., to conjunction and the nesting of existential restrictions in the C_{ij} . This set is made finite by imposing further restrictions on the generation process such as the maximal number of conjuncts and the maximal depth of existential nestings allowed. \square

Remark 4 One may also think of further partitioning the set \mathcal{E}^- of non-positive examples into a set \mathcal{E}^- of negative and a set \mathcal{E}^u of unknown examples (and use as labelling set $\{-1, 0, 1\}$, respectively, with 1 –positive, 0 – unknown, -1 – negative), as done in some other approach (see e.g. [31]). That is, an individual a is a negative example of T if $\mathcal{K} \models a:\neg T$, while a is an unknown example of T if neither $\mathcal{K} \models a:T$ nor $\mathcal{K} \models a:\neg T$ hold. In that case, usually we are looking for an exact definition of T , i.e. a hypothesis is of the stronger form $T = C$ instead.⁹ Which one to choose may depend on the application domain and on the effectiveness of the approach. We do not address this case here. \square

We conclude with the notions of *consistent*, *non-redundant*, *sound*, *complete* and *strongly complete* hypothesis h w.r.t. \mathcal{K} , which are defined as follows:

Consistency. $\mathcal{K} \cup h$ is a consistent;

Non-Redundancy. $\mathcal{K} \not\models \phi$, for all $\phi \in h$.

Soundness. $\forall a \in \mathcal{I}_{\mathcal{E}^-}, h(a) = 0$.

Completeness. $\forall a \in \mathcal{I}_{\mathcal{E}^+}, h(a) > 0$.

Strong Completeness. $\forall a \in \mathcal{I}_{\mathcal{E}^+}, h(a) = 1$.

We say that a hypothesis h *covers* (*strongly covers*) an example $e \in \mathcal{E}$ iff $\text{bed}(\mathcal{K} \cup h, e) > 0$ ($\text{bed}(\mathcal{K} \cup h, e) = 1$). Therefore, soundness states that a learned hypothesis is not allowed to cover a non-positive example, while the way (strong) completeness is stated guarantees that all positive examples are (strongly) covered.

In general a learned (*induced*) hypothesis h has to be *consistent*, *non-redundant* and *sound* w.r.t. \mathcal{K} , but not necessarily complete, but, of course, these conditions can also be relaxed.

⁹We recall that a hypothesis as in Eq. 7 does not allow us to infer negative instances of T , while $T = C$ does.

4.2 The Learning Algorithm Fuzzy OWL-Boost

We now present our real-valued boosting-based algorithm, which is based on a boosting schema (this section) applied to a fuzzy $\mathcal{EL}(\mathbf{D}^-)$ *weak learner* described in more detail in Section 4.3. Our learning method creates an *ensemble* of fuzzy GCIs (see Eq. 7): essentially, at each iteration our boosting algorithm invokes a weak learner that generates a set of fuzzy $\mathcal{EL}(\mathbf{D}^-)$ candidate GCIs that has the form $h_i = \{C_{i1} \sqsubseteq T, \dots, C_{ik_i} \sqsubseteq T\}$, called *weak hypothesis*, determining a change to the distribution of the weights associated with the examples. The weights of misclassified examples get increased so that a better classifier can be produced in the next round, indicating the harder examples to focus on. The weak hypotheses are then combined into a final hypothesis via a weighted sum of the weak hypotheses. We will rely on \mathbb{R} Real AdaBoost [63, 64] as boosting algorithm, while we will use a weak learner that is similar to FOIL- \mathcal{DL} [47, 48, 50], both of which need to be adapted to our specific setting.

Formally, consider a KB \mathcal{K} , a training set \mathcal{E} , a set of individuals \mathbf{l} with $\mathbf{l}_{\mathcal{E}} \subseteq \mathbf{l} \subseteq \mathbf{l}_{\mathcal{K}}$, and a weight distribution \mathbf{w} over \mathbf{l} . With \mathbf{u} we indicate the uniform distribution over \mathbf{l} , *i.e.* $u_a = 1/|\mathbf{l}|$ (with $a \in \mathbf{l}$). Furthermore, consider a *weak hypothesis* h_i , *i.e.* a set $h_i = \{C_{i1} \sqsubseteq T, \dots, C_{ik_i} \sqsubseteq T\}$ of fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCIs returned by the weak learner. Note that for $a \in \mathbf{l}_{\mathcal{K}}$, $\text{bed}(\mathcal{K} \cup h_i, a:T) \in [0, 1]$. Next, we transform this value into a value in $[-1, 1]$ as required by \mathbb{R} Real AdaBoost. So, let $t: [0, 1] \rightarrow [-1, 1]$ be the transformation function

$$t(x) = \begin{cases} -1 & \text{if } x = 0 \\ x & \text{else} \end{cases}$$

and let the classification prediction value $h_i(a)$ of a w.r.t. h , T and \mathcal{K} be defined as (again for ease, we omit \mathcal{K} and T)

$$h_i(a) = t(\text{bed}(\mathcal{K} \cup h_i, a:T)) \in \{-1\} \cup (0, 1] .$$

We also define the examples labelling l over \mathbf{l} in the following way: for $a \in \mathbf{l}$

$$l(a) = \begin{cases} 1 & \text{if } (a, 1) \in \mathcal{E}^+ \\ -1 & \text{else} . \end{cases}$$

Finally, for a weak hypothesis h_i , we determine also the *error* of h_i w.r.t. a distribution \mathbf{w} as

$$\epsilon(\mathbf{w}) = \sum_{a \in \mathbf{l}} w_a \cdot \delta(h_i(a), l(a)) \cdot h_i(a) ,$$

where $\delta(x, y) \in \{0, 1\}$ is defined as $(x \in \{-1\} \cup (0, 1], y \in \{-1, 1\})$

$$\delta(x, y) = \begin{cases} 1 & \text{if } x \cdot y < 0 \\ 0 & \text{else} . \end{cases}$$

Note that $\delta(h_i(a), l(a))$ determines whether there is a *disagreement* between the sign of $h_i(a)$ and $l(a)$.

Then, the FUZZY OWL-BOOST algorithm calling iteratively a weak learner is shown in Algorithm 1, which we comment briefly next.

The algorithm is similar as \mathbb{R} Real AdaBoost, except for some context dependent parts. In Step 2 we initialise the set of individuals \mathbf{l} to be considered as $\mathbf{l}_{\mathcal{K}}$. Essentially, all individuals will be weighted. The main loop (Steps 5 - 13) is similar to \mathbb{R} Real AdaBoost with the particularity that

Algorithm 1 FUZZY OWL-BOOST

Input: KB \mathcal{K} , training set \mathcal{E} , target concept name T , number of iterations n

Output: Hypothesis h as by Eq. 7.

```
1:  $h \leftarrow \emptyset$ ;
2:  $l \leftarrow l_{\mathcal{K}}$ ;
3:  $\mathbf{w}_1 \leftarrow \mathbf{u}$ ; ▷ Initialize the weight distribution over  $l$ 
4: // Main boosting loop
5: for  $i = 1$  to  $n$  do
6:    $h_i \leftarrow \text{FUZZYWEAKLEARNER}(\mathcal{K}, T, \mathcal{E}, \mathbf{w}_i)$ ; ▷ Weak learner  $h_i$ , i.e. set of axioms  $C_{ij} \sqsubseteq T$ 
7:   if  $\epsilon(\mathbf{w}_i) \geq 0.5$  then break; ▷ If weak learner error is not below 0.5 then break the loop
8:    $h_i^* \leftarrow \max_{a \in l} |h_i(a)|$ ; ▷  $h_i^*$  is the maximal value of  $h_i$  over  $l$ 
9:    $\mu_i \leftarrow \frac{1}{h_i^*} \sum_{a \in l} w_{i,a} \cdot l(a) \cdot h_i(a)$ ; ▷  $\mu_i$  is the normalized margin of  $h_i$  w.r.t.  $l$ 
10:   $\alpha_i \leftarrow \frac{1}{2h_i^*} \cdot \ln \frac{1+\mu_i}{1-\mu_i}$ ; ▷  $\alpha_i$  is the weight of classifier  $h_i$  in the ensemble
11:  for all  $a \in l$  do ▷ Update the weight distribution
12:     $w_{i+1,a} \leftarrow w_{i+1,a} \cdot \left( \frac{1 - (\mu_i \cdot l(a) \cdot h_i(a)) / h_i^*}{1 - \mu_i^2} \right)$ ;
13:   $h \leftarrow h \cup \{C_{ij} \sqsubseteq WL_i \mid C_{ij} \sqsubseteq T \in h_i, WL_i \text{ new}\}$  ▷ Update hypothesis according to Eq. 7.
14: // Build now the final classifier ensemble
15:  $\phi_T \leftarrow \alpha_1 \cdot WL_1 + \dots + \alpha_n \cdot WL_n \sqsubseteq T$ ;
16:  $h \leftarrow h \cup \{\phi_T\}$ ;
17: return  $h$ ;
```

Step 6 we invoke our weak learner that is assumed to return a set $h_i = \{C_{i1} \sqsubseteq T, \dots, C_{ik_i} \sqsubseteq T\}$ of fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCIs. In Step 7 we have a case that causes a break of the main loop. In fact, an implicit condition of boosting is that the error of a weak learner should be below 0.5. That is, the weak hypothesis should be better than random guess.

In Step 13 we update the hypothesis h with the weak hypothesis, while in Steps 15 - 16 we build the final classifier ensemble and add it to the hypothesis.

4.3 The Weak Learner wFoil- \mathcal{DL}

We next describe the weak learner we employ here. Specifically, we will use a FOIL- \mathcal{DL} [47, 48, 50] like weak learner, which however needs to be adapted to our specific setting (see Algorithm 2). In general terms the weak learning algorithm, called wFOIL- \mathcal{DL} , follows a so-called *sequential covering* learning approach. That is, one carries on inducing GCIs until all positive examples are covered or nothing new can be learned. When an axiom is induced (see Step 4 in Algorithm 2), the positive examples still to be covered are updated (Step 9 and 10).

In order to induce an axiom (Step 4), LEARN-ONE-AXIOM is invoked (see Algorithm 3), which in general terms operates as follows:

1. start from concept \top ;
2. apply a refinement operator to find more specific fuzzy $\mathcal{EL}(\mathbf{D}^-)$ concept description candidates;
3. exploit a scoring function to choose the best candidate;
4. re-apply the refinement operator until a good candidate is found;
5. iterate the whole procedure until a satisfactory coverage of the positive examples is achieved.

Algorithm 2 wFOIL- \mathcal{DL} (Weak Learner)

Input: KB \mathcal{K} , target concept name T , training set \mathcal{E} , weight distribution \mathbf{w} , confidence threshold $\theta \in [0, 1]$, non-positive coverage percentage $\eta \in [0, 1]$

Output: A weak hypothesis, *i.e.* a set $h = \{C_1 \sqsubseteq T, \dots, C_k \sqsubseteq T\}$ of fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCIs

- 1: $h \leftarrow \emptyset, Pos \leftarrow \mathcal{E}^+, \phi \leftarrow \top \sqsubseteq T$;
- 2: //Loop until no improvement
- 3: **while** ($Pos \neq \emptyset$) **and** ($\phi \neq \text{null}$) **do**
- 4: $\phi \leftarrow \text{LEARN-ONE-AXIOM}(\mathcal{K}, T, \mathcal{E}, \mathbf{w}_i, Pos, \theta, \eta)$; // Learn one fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCI of the form $C \sqsubseteq T$
- 5: **if** $\phi \in h$ **then** // axiom already learned
- 6: $\phi \leftarrow \text{null}$;
- 7: **if** $\phi \neq \text{null}$ **then**
- 8: $h \leftarrow h \cup \{\phi\}$; // Update weak hypothesis
- 9: $Pos_\phi \leftarrow \{\langle a, 1 \rangle \in \mathcal{E}^+ \mid \text{bed}(\mathcal{K} \cup \{\phi\}, T(a)) > 0\}$; // Positives covered by ϕ
- 10: $Pos \leftarrow Pos \setminus Pos_\phi$; // Update positives still to be covered
- 11: **return** h ;

Table 1: Downward Refinement Operator.

$$\rho(C) = \begin{cases} \mathbf{A}_{\mathcal{K}} \cup \{\exists r. \top \mid r \in \mathbf{R}_{\mathcal{K}}\} \cup \{\exists s. d \mid s \in \mathbf{S}_{\mathcal{K}}, d \in \mathcal{D}\} \cup & \text{if } C = \top \\ \{\exists s. =_b, \mid s \in \mathbf{B}_{\mathcal{K}}, b \in \{\text{true}, \text{false}\}\} & \text{if } C = A \\ \{A' \mid A' \in \mathbf{A}_{\mathcal{K}}, \mathcal{K} \models A' \sqsubseteq A\} \cup \{A \sqcap A'' \mid A'' \in \rho(\top)\} & \text{if } C = \exists r. D, r \in \mathbf{R}_{\mathcal{K}} \\ \{\exists r. D' \mid D' \in \rho(D)\} \cup \{\exists r. D \sqcap D'' \mid D'' \in \rho(\top)\} & \text{if } C = \exists s. d, s \in \mathbf{S}_{\mathcal{K}}, d \in \mathcal{D} \\ \{\exists s. d \sqcap D \mid D \in \rho(\top)\} & \text{if } C = \exists s. =_b, s \in \mathbf{B}_{\mathcal{K}}, b \in \{\text{true}, \text{false}\} \\ \{\exists s. =_b \sqcap D \mid D \in \rho(\top)\} & \text{if } C = C_1 \sqcap \dots \sqcap C_n \\ \{C_1 \sqcap \dots \sqcap C'_i \sqcap \dots \sqcap C_n \mid i = 1, \dots, n, C'_i \in \rho(C_i)\} & \text{if } C = C_1 \sqcap \dots \sqcap C_n \end{cases}$$

We briefly detail the steps of LEARN-ONE-AXIOM.

Computing fuzzy datatypes. For a numerical datatype s , we consider *equal width triangular partitions* of values $V_s = \{v \mid \mathcal{K} \models a: \exists s. =_v\}$ into a finite number of fuzzy sets (3, 5 or 7 sets), which is identical to [47, 50, 85] (see, *e.g.* Fig. 3). However, we additionally, allow also the use of the C-means fuzzy clustering algorithm over V_s , where the fuzzy membership function is a triangular function build around the centroid of a cluster. Note that C-means has not been considered in [47, 50, 85].¹⁰

The refinement operator. The refinement operator we employ is the same as in [47, 48, 54, 85] except that now we add the management of boolean values as well. Essentially, the refinement operator takes as input a concept C and generates new, more specific concept description candidates D (*i.e.*, $\mathcal{K} \models D \sqsubseteq C$). For the sake of completeness, we recap the refinement operator here. Let \mathcal{K} be a knowledge base, $\mathbf{A}_{\mathcal{K}}$ be the set of all atomic concepts in \mathcal{K} , $\mathbf{R}_{\mathcal{K}}$ the set of all object properties in \mathcal{K} , $\mathbf{S}_{\mathcal{K}}$ the set of all numeric datatype properties in \mathcal{K} , $\mathbf{B}_{\mathcal{K}}$ the set of all boolean datatype properties in \mathcal{K} and \mathcal{D} a set of (fuzzy) datatypes. The refinement operator ρ is shown in Table 1.

The scoring function. The scoring function we use to assign a score to each candidate hypothesis is essentially a weighted *gain* function, similar to the one employed in [47, 48, 54, 85] and implements an information-theoretic criterion for selecting the best candidate at each refinement

¹⁰Specifically, C-means has not been considered so far in fuzzy GCI learning.

step. Specifically, given a fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCI ϕ of the form $C \sqsubseteq T$ chosen at the previous step, a KB \mathcal{K} , a set of individuals \mathbf{l} , a weight distribution \mathbf{w} over \mathbf{l} , a set of positive examples Pos still to be covered and a candidate fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCI ϕ' of the form $C' \sqsubseteq T$, then

$$gain(\phi', \phi, \mathbf{w}, \mathbf{l}, Pos) = p * (\log_2(cf(\phi', \mathbf{w}, \mathbf{l}, Pos)) - \log_2(cf(\phi, \mathbf{w}, \mathbf{l}, Pos))) , \quad (8)$$

where $p = |C' \cap C|_{\mathcal{K}}^{\mathbf{w}, \mathbf{l}, Pos}$ is the weighted cardinality of positive examples in Pos covered by ϕ' that are still covered by ϕ , and

$$cf(D \sqsubseteq T, \mathbf{w}, \mathbf{l}, Pos) = \frac{|D|_{\mathcal{K}}^{\mathbf{w}, \mathbf{l}, Pos}}{|D|_{\mathcal{K}}^{\mathbf{w}, \mathbf{l}}} . \quad (9)$$

Please note that in Eq. 9, about the confidence degree of $D \sqsubseteq T$, the numerator is calculated w.r.t. the positive examples still to be covered, *i.e.* all instances of T that are in Pos and are instances of D . In this way, LEARN-ONE-AXIOM is somewhat guided towards positives not yet covered by the weak learner learned so far by wFOIL-DL. Note also that the gain is positive if the confidence degree increases.

Stop criterion. LEARN-ONE-AXIOM stops when the confidence degree is above a given threshold $\theta \in [0, 1]$, or no GCI can be found that does not cover any negative example (in \mathcal{E}^-) above a given percentage.

The Learn-One-Axiom algorithm. The LEARN-ONE-AXIOM algorithm is defined in Algorithm 3, which we comment briefly as next. Steps 1 - 3 are simple initialisation steps. Steps 5 - 21 are the main loop from which we may exit in case there is no improvement (Step 16), and the confidence degree of the so far determined GCI is above a given threshold or it does not cover any negative example above a given percentage (Step 18). Note that the latter case guarantees soundness of the weak learner if this percentage is set to 0. In Step 8 we determine all new refinements, which then are scored in Steps 10 - 15 in order to determine the one with the best gain. At the end of the algorithm, once we exit from the main loop, the best found GCI is returned (Step 22).

Remark 5 *As for FOIL-DL (and pFOIL-DL), the weak learner wFOIL-DL also allows to use a backtracking mechanism (Step 19), which, for ease of presentation, we omit to include. The mechanism is exactly the same as for the pFOIL-DL-learnOneAxiom described in [85, Algorithm 3]. Essentially, a stack of top-k refinements is maintained, ranked in decreasing order of the confidence degree from which we pop the next best refinement (if the stack is not empty) in case no improvement has occurred. C_{best} becomes the popped-up refinement.*

5 Evaluation

We have implemented the algorithm within the *FuzzyDL-Learner*¹¹ system and evaluated it over a set of (crisp) OWL ontologies. All the data and implementation can be downloaded from the FuzzyDL-Learner home page.

5.1 Setup

A number of OWL ontologies from different domains have been selected as illustrated in Tables 2 and 3. A succinct description of them is provided in Appendix A. Note that the ontologies in

¹¹<http://www.umbertostraccia.it/cs/software/FuzzyDL-Learner/>.

Algorithm 3 LEARN-ONE-AXIOM

Input: KB \mathcal{K} , target concept name T , training set \mathcal{E} , weight distribution \mathbf{w} , Pos set of positive examples still to be covered, confidence threshold $\theta \in [0, 1]$, non-positive coverage percentage $\eta \in [0, 1]$

Output: A fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCI of the form $C \sqsubseteq T$

```
1:  $l \leftarrow l_{\mathcal{K}}$ ;
2:  $C \leftarrow \top$ ; ▷ Start from  $\top$ 
3:  $\phi \leftarrow C \sqsubseteq T$ ;
4: //Loop until no improvement
5: while  $C \neq \text{null}$  do
6:    $C_{best} \leftarrow C$ ;
7:    $maxgain \leftarrow 0$ ;
8:    $\mathcal{C} \leftarrow \rho(C)$ ; ▷ Compute all refinements of  $C$ 
9:   // Compute the score of the refinements and select the best one
10:  for all  $C' \in \mathcal{C}$  do
11:     $\phi' \leftarrow C' \sqsubseteq T$ ;
12:     $gain \leftarrow gain(\phi', \phi, \mathbf{w}, l, Pos)$ ;
13:    if ( $gain > maxgain$ ) and ( $cf(\phi', \mathbf{w}, l, Pos) > cf(\phi, \mathbf{w}, l, Pos)$ ) then
14:       $maxgain \leftarrow gain$ ;
15:       $C_{best} \leftarrow C'$ ;
16:  if  $C_{best} = C$  then ▷ No improvement
17:    //Stop if confidence degree above threshold or no negative coverage below threshold
18:    if ( $cf(C_{best} \sqsubseteq T, l) \geq \theta$ ) and ( $\frac{[C_{best}]_{\mathcal{K}^-}^{\mathcal{E}^-}}{|l_{\mathcal{E}^-}|} \leq \eta$ ) then break;
19:    // Manage backtrack here, if foreseen
20:     $C \leftarrow C_{best}$ ;
21:     $\phi \leftarrow C \sqsubseteq T$ ;
22: return  $\phi$ ;
```

Table 2: Facts about the ontologies of the experiment.

ontology	DL	class.	obj. prop.	data. prop.	ind.	target T	pos	neg	dth/cj/ η
FamilyTree	$SROLF(\mathcal{D})$	22	52	6	368	Uncle	46	156	1/5/0
Hotel	$ALCCOF(\mathcal{D})$	89	3	1	88	Good_Hotel	12	11	1/5/0
Moral	ALC	46	0	0	202	ToLearn_Guilty	102	100	1/5/0
SemanticBible (NTN)	$SHOLN(\mathcal{D})$	51	29	9	723	ToLearn_Woman	46	3	1/5/0
UBA	$SHI(\mathcal{D})$	44	26	8	1268	Good_Researcher	22	113	1/5/0
WineOnto	$SHI(\mathcal{D})$	178	15	7	138	ToLearn_DryWine	15	-	1/5/0
Pair50	ALC	3	6	0	311	ToLearn	20	29	2/5/0
Straight	ALC	3	6	0	347	ToLearn	4	50	3/5/1.0
Lymphography	ALC	50	0	0	148	ToLearn	81	67	1/5/1.0
Mammographic	$ALC(\mathcal{D})$	20	3	2	975	ToLearn	445	516	3/5/1.0
Pyrimidine	$ALC(\mathcal{D})$	2	0	27	74	ToLearn	20	20	1/5/1.0
Suramin	$ALC(\mathcal{D})$	47	3	1	2979	ToLearn	7	10	3/5/1.0

Table 3: Datasets considered from the UCI ML Repository.

dataset	instances	attributes	target T	pos	dth/cj/ η
Iris	151	4	Iris-setosa	51	1/5/1.0
			Iris-versicolor	50	
			Iris-virginica	50	
Wine	178	13	1, 2, 3	59, 71, 48	1/5/1.0
Wine Quality	4898	12	GoodRedWine	18	1/5/1.0

Table 3 are not available as OWL 2 ontologies but only as csv format. Therefore, we have translated them from the csv format according to the procedure shown in Appendix B. Furthermore, note that the ontologies in Table 3 are taken from the well-known *UC Irvine Machine Learning Repository* [23]. While evaluating ontology-based learning algorithms is untypical on numerical datatype properties¹², we believe it is interesting to do so as an important ingredient of our algorithm is the use of fuzzy concrete datatype properties.

For each ontology \mathcal{K} a meaningful target concept has been selected such that the conditions of the learning problem are satisfied. We report also the DL the ontology refers to, the number of concept names, object properties, datatype properties and individuals in the ontology. We also report the maximal nesting depth (dth.), maximal number of conjuncts (cj.) and maximal percentage of false positives (η) during the learning phase. The number n of iterations of FUZZY OWL-BOOST is set to 10.¹³ We did not consider backtracking. Nevertheless, all configuration parameters for each run are available from the downloadable data.

We will consider the following effectiveness measures (see also [85] for similar measures), which we report here for clarity to avoid ambiguity. Specifically, consider the classifier ensemble h returned by FUZZY OWL-BOOST and let us assume to have added it to the KB \mathcal{K} . Then we consider the following fuzzy measures, were their well-known crisp variants [6] (in the denotation we omit the f subscript) are obtained by replacing in the equations below the cardinality function $|\cdot|_{\mathcal{K}}$ (see Eq. 3) with the crisp cardinality function $[\cdot]_{\mathcal{K}}$.

Fuzzy True Positives: denoted TP_f , is defined as

$$TP_f = |T|_{\mathcal{K}}^{\dagger} , \quad (10)$$

¹²To the best of our knowledge, we are unaware of any evaluation of ontology-based methods on those data sets.

¹³We tried also for $n > 10$ and did not notice positive effects. In fact, at some point the weak learner is unable to learn new rules given the weight distribution.

Fuzzy False Positives: denoted FP_f , is defined as

$$FP_f = |T|_{\mathcal{K}}^{\varepsilon^-} , \quad (11)$$

Fuzzy True Non-Positive: denoted TNP_f , is defined as

$$TNP_f = |\mathcal{I}_{\varepsilon^-}| - FP , \quad (12)$$

Fuzzy False Non-Positive: denoted FNP_f , is defined as

$$FNP_f = |\mathcal{I}_{\varepsilon^+}| - TP , \quad (13)$$

Fuzzy Precision: denoted P_f , is defined as

$$P_f = \frac{TP_f}{|T|_{\mathcal{K}}^{\varepsilon}} , \quad (14)$$

Fuzzy Recall: denoted R_f , is defined as

$$R_f = \frac{TP_f}{|\mathcal{I}_{\varepsilon^+}|} , \quad (15)$$

Fuzzy F1-score: denoted $F1_f$, is defined as

$$F1_f = 2 \cdot \frac{P_f \cdot R_f}{P_f + R_f} ,$$

Mean Squared Error: denoted MSE , is defined as

$$MSE = \frac{1}{|\mathcal{I}_{\varepsilon}|} \cdot \sum_{a \in \mathcal{I}_{\varepsilon}} (h(a) - l(a))^2 .$$

Concerning other parameter settings, we

- varied the number of fuzzy sets (3, 5 or 7). For C-means, we fixed the hyper-parameter to the default $m = 2$, the threshold to $\epsilon = 0.05$ and the number of maximum iterations to 100; and
- varied the confidence threshold $\theta \in \{0.34, 0.64, 0.94, 1.0\}$.

Therefore, we considered a total of 12 different parameter configurations.

For each parameter configuration, a stratified k -fold cross validation design¹⁴ was adopted (specifically, $k = 5$) to determine the average of the above described performance indices. For each measure, the (macro) average value over the various folds has been considered. In all tests, we have that $\mathcal{I}_{\varepsilon} = \mathcal{I}_{\mathcal{K}}$ and that there is at least one positive example in each fold, while the other examples of a fold have been randomly been selected. For each fold, all assertions involving testing examples have been removed from a given ontology, thus restricting the training phase to training examples only. We considered also the extreme case in which the whole set \mathcal{E} is used for both training and testing. This case has been considered for those ontologies with few

¹⁴Stratification means here that each fold contains roughly the same proportions of positive and non-positive instances of the target class.

positive examples for which k -fold cross validation is not meaningful and also for the task aiming at “explaining” the target w.r.t. the given data set. This case is indicated with \star in Table 4.

As baseline, we considered an improved version of FOIL- \mathcal{DL} w.r.t. the one published in [47, 48, 50]. Roughly, FOIL- \mathcal{DL} , as wFOIL- \mathcal{DL} , learns iteratively rules. At each iteration i , FOIL- \mathcal{DL} learns one fuzzy $\mathcal{EL}(\mathbf{D}^-)$ GCI of the form $\langle C_i \sqsubseteq T, d_i \rangle$ by invoking a similar procedure as LEARN-ONE-AXIOM, where however

- d_i is the confidence degree of $C_i \sqsubseteq T$; and
- the weight distribution \mathbf{w} is roughly as follows: if a positive instance a has already been covered by the rules learned so far, then $w_a = 0$ (this is the same as to say to remove the covered positive instances from the next iteration). The weight of the other instances is determined according to a *uniform* distribution.

At the end, the final hypothesis of FOIL- \mathcal{DL} is of the form (cf. Eq. 7)

$$\begin{aligned} &\langle C_1 \sqsubseteq T, d_1 \rangle \\ &\quad \vdots \\ &\langle C_n \sqsubseteq T, d_n \rangle . \end{aligned} \tag{16}$$

Therefore, there is a notable difference among FUZZY OWL-BOOST and FOIL- \mathcal{DL} in (i) the way the instance distribution is set up at each iteration; (ii) how the weight of each rule is determined (the α_i in FUZZY OWL-BOOST versus the confidence d_i in FOIL- \mathcal{DL}); and (iii) how the final hypothesis is build (linear combination in FUZZY OWL-BOOST versus ‘max’ aggregation in FOIL- \mathcal{DL}).

In the result Table 4, for a given KB \mathcal{K} , a given algorithm (FUZZY OWL-BOOST or FOIL- \mathcal{DL}) and a given clustering method (uniform u or C-means c), we report only the effectiveness measures for the configuration (θ, fs) ¹⁵ with the highest score of

$$fF1F1 = F1_f \cdot F1 , \tag{17}$$

i.e. a compromise (Pareto optimal solution) among fuzzy $F1$ and crisp $F1$, as, more often than not, the best fuzzy $F1$ and best crisp $F1$ values do not relate to the same configuration.¹⁶

Example 5.1 *We provide here an example of learned rule set (in Manchester OWL syntax) via FUZZY OWL-BOOST applied to the Wine dataset (see Table 3) and target class 2 (considering the best run).*

```
# Weak Learner WL1
(Alcohol some Alcohol_VL) and (Hue some Hue_H) and (MalidAcid some MalidAcid_VVL) SubClassOf WL1
(Ash some Ash_L) and (ColorIntensity some ColorIntensity_VVL) SubClassOf WL1
(ColorIntensity some ColorIntensity_VVL) and (Proline some Proline_L) SubClassOf WL1
(ColorIntensity some ColorIntensity_VVL) and (Proline some Proline_VVL) SubClassOf WL1
(Magnesium some Magnesium_VVL) and (NonFlavonoidsPhenols some NonFlavonoidsPhenols_VVH) and (Proanthocyanins some Proanthocyanins_F) SubClassOf WL1
(MalidAcid some MalidAcid_VVL) and (Proline some Proline_VVL) SubClassOf WL1

# Weak Learner WL2
(Alcohol some Alcohol_VVL) and (ColorIntensity some ColorIntensity_VVL) SubClassOf WL2
(Ash some Ash_VVL) and (Flavonoids some Flavonoids_L) SubClassOf WL2
(ColorIntensity some ColorIntensity_VVL) and (Proline some Proline_VVL) SubClassOf WL2

# Weak Learner WL3
(ColorIntensity some ColorIntensity_VVL) SubClassOf WL3

# Real Adaboost aggregation
1.199 * WL1 + 0.544 * WL2 + 0.272 * WL3 SubClassOf 2

# Fuzzy datatypes (C-Means)
Alcohol_VVL left-shoulder 11.81 12.29
Alcohol_VL triangular 11.79 12.27 12.58
```

¹⁵Recall that $\theta \in \{0.34, 0.64, 0.94, 1.0\}$, $fs \in \{3, 5, 7\}$.

¹⁶In case of a tie, we adopt the following priorities: lowest θ and then lowest number of partitions.

Table 4: Results table.

Dataset, Target	Algorithm	Clustering	theta	fs	MSE	Fuzzy F1	F1	FuzzyF1 * F1	% Improvement
Iris (setosa)	FOIL-DL	u	0.34	3	0.018	0.893	1.000	0.893	
	Fuzzy OWL-Boost	u	0.34	3	0.000	1.000	1.000	1.000	11.93%
Iris (versicolor)	FOIL-DL	c	0.94	7	0.090	0.710	0.888	0.630	
	Fuzzy OWL-Boost	c	0.64	3	0.021	0.929	0.860	0.799	26.68%
Iris (virginica)	FOIL-DL	u	0.94	5	0.048	0.851	0.891	0.758	
	Fuzzy OWL-Boost	u	0.64	5	0.036	0.922	0.891	0.821	8.30%
Wine (1)	FOIL-DL	u	0.94	5	0.076	0.742	0.915	0.679	
	Fuzzy OWL-Boost	u	0.94	3	0.049	0.897	0.907	0.813	19.79%
Wine (2)	FOIL-DL	c	0.94	7	0.074	0.837	0.885	0.741	
	Fuzzy OWL-Boost	c	0.94	7	0.083	0.844	0.910	0.769	3.75%
Wine (3)	FOIL-DL	c	0.64	7	0.033	0.868	0.885	0.768	
	Fuzzy OWL-Boost	u	0.94	7	0.017	0.955	0.932	0.890	15.86%
Wine Quality	FOIL-DL	c	0.94	7	0.003	0.034	0.053	0.002	
	Fuzzy OWL-Boost	c	0.04	5	0.011	0.047	0.065	0.003	68.49%
FamilyTree	FOIL-DL	u/c	0.94	5, 3	0.019	0.929	0.929	0.863	
	Fuzzy OWL-Boost	u/c	0.94	3	0.018	0.924	0.924	0.853	-1.20%
Hotel (*)	FOIL-DL	c	0.94	7	0.007	0.942	1.000	0.942	
	Fuzzy OWL-Boost	c	0.34	7	0.005	0.968	1.000	0.968	2.76%
Moral	FOIL-DL	-	0.34	3	0.000	1.000	1.000	1.000	
	Fuzzy OWL-Boost	-	0.34	3	0.000	1.000	1.000	1.000	0.00%
SemanticBible (NTN)	FOIL-DL	u/c	0.34	3	0.046	0.527	0.548	0.289	
	Fuzzy OWL-Boost	u/c	0.94	3	0.040	0.541	0.541	0.293	1.46%
UBA (*)	FOIL-DL	c	0.34	3	0.001	0.972	0.977	0.950	
	Fuzzy OWL-Boost	c	1.00	3	0.000	1.000	1.000	1.000	5.30%
Pair50 (*)	FOIL-DL	-	0.34	3	0.000	1.000	1.000	1.000	
	Fuzzy OWL-Boost	-	0.34	3	0.000	1.000	1.000	1.000	0.00%
Straight (*)	FOIL-DL	-	0.34	3	0.007	0.400	0.571	0.228	
	Fuzzy OWL-Boost	-	0.34	3	0.021	0.514	0.500	0.257	12.52%
WineOnto (*)	FOIL-DL	-	0.34	3	0.000	1.000	1.000	1.000	
	Fuzzy OWL-Boost	-	0.34	3	0.000	1.000	1.000	1.000	0.00%
Lymphography	FOIL-DL	-	0.34	3	0.159	0.845	0.855	0.722	
	Fuzzy OWL-Boost	-	0.64	3	0.160	0.861	0.870	0.749	3.71%
Mammographic	FOIL-DL	u	0.34	3	0.222	0.593	0.710	0.421	
	Fuzzy OWL-Boost	c	0.64	7	0.199	0.720	0.774	0.557	32.26%
Pyrimidine (*)	FOIL-DL	u	1.00	7	0.050	0.889	0.947	0.842	
	Fuzzy OWL-Boost	c	0.34	7	0.027	0.952	0.930	0.885	5.16%
Suramin (*)	FOIL-DL	c	0.34	3	0.002	0.291	0.583	0.170	
	Fuzzy OWL-Boost	u/c	0.34	3	0.003	0.583	0.583	0.340	100.34%

Hue_H triangular 1.00 1.08 1.21
MalidAcid_VWL left-shoulder 1.27 1.67
Ash_VWL left-shoulder 1.88 2.11
Ash_L triangular 2.12 2.25 2.35
ColorIntensity_VWL left-shoulder 2.66 3.64
Proline_L triangular 532.73 657.29 811.52
Proline_VWL left-shoulder 424.58 532.73
Magnesium_VWL left-shoulder 85.39 89.97
NonFlavonoidsPhenols_VVH right-shoulder 0.50 0.56
Proanthocyanins_F triangular 1.40 1.64 1.94
Flavonoids_L triangular 1.35 1.76 2.21

□

5.2 Discussion

We now discuss the results in Table 4. We report in red the percent improvement of FUZZY OWL-BOOST, relative to the measure $fF1F1$ (see Eq. 17), over our baseline FOIL- \mathcal{DL} .

Uniform vs. C-Means fuzzy datatype construction To start with, without going to much into it as it is not the main of this work, not surprisingly C-means behaves better than the ‘uniform’ (u) approach (14 wins vs. 9).¹⁷ Moreover, concerning the number of fuzzy set partitions, there is no clear indication about which choice between 3 or 7 partitions is the better way to go. The choice seems dependent on the dataset. Apparently, if 3 partitions are not enough, then one may go for 7 as likely the dataset may require a more fine grained approach. Of course, the results of C-means may further be improved by optimising its parameters. Nevertheless, the uniform approach performed surprisingly well, despite its simplicity.

¹⁷We do not count ties.




Figure 5: Improvement of FUZZY OWL-BOOST over FOIL- \mathcal{DL} according to the $fF1F1$ (Eq. 17) measure.

A more in depth investigation will be subject of future work in which we will consider some more options for fuzzy set construction and focus on its impact on the overall effectiveness.

Fuzzy OWL-Boost vs. Foil- \mathcal{DL} It appears evident from the results in Table 4 (see also Figure 5) that FUZZY OWL-BOOST performs generally better than FOIL- \mathcal{DL} (15 wins vs. 1), with 3 ties and 1 loss. Concerning the ties, note that the value of $fF1F1$ is 1.0 for both FUZZY OWL-BOOST and FOIL- \mathcal{DL} and, thus, there was no margin for improvement for FUZZY OWL-BOOST. The only loss was for the FamilyTree dataset, though, the difference is small (-1.20%).

The average improvement of FUZZY OWL-BOOST over all runs is 16.69%. This is essentially due to the average improvement w.r.t. the fuzzy F1 measure, which is 15.16% (15 wins vs. 1, 3 ties), while the average improvement for the (crisp) F1 measure is marginal 1.21% (6 wins vs. 7, 6 ties).

Overall, note also that the average MSE is low (0.041) with a slightly advantage for FUZZY OWL-BOOST over FOIL- \mathcal{DL} (0.036 vs. 0.045). Nevertheless, there are two outlier datasets (Lymphography and Mammographic) for which there is quite some room for improvement of the MSE .

Last but not least, let us mention that both FUZZY OWL-BOOST and FOIL- \mathcal{DL} do definitely not behave well on the WineQuality dataset, which will be the subject of further investigation.¹⁸

6 Conclusions & Future Work

In this work, we addressed the problem of automatically learning fuzzy concept inclusion axioms from OWL 2 ontologies. That is, given a target class T of an OWL ontology, we address the problem of inducing a fuzzy $\mathcal{ELW}(\mathbf{D})$ concept inclusion axioms that describe sufficient conditions for being an individual instance of T . In particular, we have adapted the Real AdaBoost [64] boosting algorithm to the fuzzy OWL case, by presenting the FUZZY OWL-BOOST algorithm. The main features of our algorithm are essentially the fact that (i) it generates a set of fuzzy $\mathcal{EL}(\mathbf{D}^-)$ inclusion axioms, which are the weak hypothesis, possibly including fuzzy concepts and

¹⁸However, a run on this dataset requires ca. one week of computation on our hardware (Linux OS, with 16GB RAM and Intel Core i9-9900K CPU @ 3.60GHz).

fuzzy concrete domains; *(ii)* combines them via a weighted sum; and *(iii)* all generated fuzzy concept inclusion axioms can be encoded as *Fuzzy OWL 2* axioms.

We have also conducted an extensive evaluation, comparing FUZZY OWL-BOOST with FOIL- \mathcal{DL} . Our evaluation shows that FUZZY OWL-BOOST is generally better than FOIL- \mathcal{DL} in terms of *fF1F1* effectiveness (+16.69% average) over the tested datasets, and that the improvement mainly concerns the fuzzy F1 measure, while effectiveness remains essentially similar for (crisp) F1. Also, the C-means clustering method prevails over the uniform clustering method to build the fuzzy data types. Let us also note that both FUZZY OWL-BOOST (as well as FOIL- \mathcal{DL}) generates easy human interpretable hypotheses (see *e.g.* Example 5.1).

Last but not least, let us mention that in a previous version of this work, we also considered the use of the LEARN-ONE-AXIOM algorithm only (see Algorithm 3) as weak learner in place of wFOIL- \mathcal{DL} and the softmax function to normalise the weights α_i in the weighted sum construct. However, the results were not really encouraging (*i.e.* slightly worse) w.r.t. FOIL- \mathcal{DL} .

Concerning future work, besides investigating about other learning methods, we envisage various aspects worth to be investigated in more detail: *(i)* we would like to make an in depth investigation about the impact of clustering methods for building fuzzy datatypes on the overall effectiveness by considering various alternatives as well, as proposed recently in a Fuzzy Sets and Systems special issue on fuzzy clustering [1]. Moreover, we would like to cover more OWL datatypes than those considered here so far (numerical and boolean) such as strings, dates, etc. possibly in combination with some sub-atomic classical machine learning methods (see, *e.g.* [81]); *(ii)* another aspect may concern the investigation of the impact of choosing various fuzzy semantics during the learning phase; *(iii)* last but not least, we would like to investigate the computational aspect: so far, for some ontologies, a learning run may take even a week (on the resource at our disposal¹⁹). We would like to investigate both parallelisation methods as well as to investigate about the impact, in terms of effectiveness, of efficient, logically sound, but not necessarily complete, reasoning algorithms, such as *structural* DL algorithms.

Acknowledgment

This research was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. This work has also been partially supported by the H2020 DeepHealth Project (GA No. 825111).

A Brief Description of the Datasets

A.1 OWL Ontologies Description

Find below a brief description about the OWL ontologies in Table 2 used in our experiments. Some ontology descriptions can also found in [86].²⁰

FamilyTree. This is a simple family relationships ontology and associated instances. The description is of the family of Robert Stevens and the intention is to use the minimal of asserted relationships and the maximum of inference. To do this, role chains, nominals and properties hierarchies have been used. The target is to identify sufficient conditions for being an uncle.

¹⁹See footnote 18.

²⁰See also, <https://github.com/SmartDataAnalytics/SML-Bench>

Hotel. This ontology describes the meaningful entities of a city. Instances are hotels located in the town Pisa and ratings have been gathered from Trip Advisor.²¹ The target is to identify sufficient conditions for being a good hotel, which has been identified as a hotel having a rating above 4.

Moral. This ontology is about meaningful entities involved in the description of guiltiness within a moral theory of blame scenario. The target is to learn sufficient conditions to be guilty.

SemanticBible (NTN). *New Testament Names* (NTN) is an ontology describing each named thing in the New Testament, about 600 names in all. Each named thing (an entity) is categorized according to its class, including God, Jesus, individual men and women, groups of people, and locations. These entities are related to each other by properties that interconnect the entities into a web of information.²² The target is to learn sufficient conditions to be a woman.

UBA. This is a well-known university ontology for benchmark tests describing meaningful entities within a university (*e.g.* universities, departments and the activities that occur at them).²³ The target here is to determine sufficient conditions to be a good researcher.

WineOnto. This is an ontology about Italian, French and German red and white wines involving the description of, among others, their chemical properties. The target here is to determine sufficient conditions to be a dry wine.

Pair50. This ontology is about a poker game and the target is to determine whether a player has a pair at hand.

Straight. This ontology is about a poker game, as the one for Pair50, but the target is now to determine whether one has a straight at hand.

Lymphography. This ontology is about lymphography patient data and the target is the prediction of a diagnosis class based on the lymphography patient data [86].

Mammographic. This ontology is about mammography screening data and the target is the prediction of breast cancer severity based on the screening data [86].

Pyrimidine. This ontology is about pyrimidine data, the target is the prediction of the inhibition activity of pyrimidines and the DHFR enzyme [86].

Suramin. This ontology is about the description of chemical compounds and the target is to find a predictive description of suramin analogues for cancer treatment.

A.2 UCI ML Data Sets

The data sets in Table 3 have been taken from the well-known *UC Irvine Machine Learning Repository* [23]. A brief description of the selected data is given below.

Iris. The data set contains 3 classes of 50 instances each, where each class refers to a type of iris plant. The attributes are: sepal length in cm, sepal width in cm, petal length in cm and petal width in cm. The target classes are: Iris Setosa, Iris Versicolour and Iris Virginica.

²¹<http://www.tripadvisor.com>

²²<http://semanticbible.com/ntn/ntn-overview.html>

²³<http://swat.cse.lehigh.edu/projects/lubm/>

Wine. These data are the results of a chemical analysis of wines grown in the same region in Italy but derived from three different cultivars. The analysis determined the quantities of 13 constituents found in each of the three types of wines. The attributes are alcohol, malic acid, ash, alcalinity of ash, magnesium, total phenols, flavonoids, nonflavonoid phenols, proanthocyanins, color intensity, hue, OD280/OD315 of diluted wines and proline. The target classes are the three wines 1, 2 and 3.

Wine Quality. The data set is related to red and white variants of the Portuguese “Vinho Verde” wine. The attributes are: fixed acidity, volatile acidity, citric acid, residual sugar, chlorides, free sulfur dioxide, total sulfur dioxide, density, pH, sulphates, alcohol and quality (score between 0 and 10). The target class is GoodRedWine, which is defined as red wines having quality score above 7.

B UCI ML Conversion Algorithm

We considered the well-known *UC Irvine Machine Learning Repository* [23] from which selected some popular datasets with numerical attributes as shown in Table 3. As anticipated, as the datasets in Table 3 are not available as OWL 2 ontologies, we have translated them from a csv format into an OWL 2 ontology in a simple way that we describe next. The method is quite general and can be applied to any other dataset with similar specifications and a dedicated procedure is available within our implemented learner for future evaluations.

Consider a dataset D with (functional) attributes s_1, \dots, s_n of type t_1, \dots, t_n . Each data record r is of the form $\langle v_1, \dots, v_n, T \rangle$, where v_i is the value of attribute s_i of type t_i , while T is the target class name for record r . For instance, for the iris dataset we have attributes

`sepal.length, sepal.width, petal.length, petal.width`

of type

`double, double, double, double`

and the first record r is

`<5.1, 3.5, 1.4, 0.2, Iris – setosa> .`

The knowledge base \mathcal{K}_D built to describe the data is as follows. Let T_D be the set of all target class names T occurring in D . The set of GCIs in \mathcal{K}_D is

$$\begin{array}{l} T \quad \sqsubseteq \quad \text{class} \quad (T \in T_D) \\ \text{class} \quad \sqsubseteq \quad \exists s_i.t_i \quad (i = 1..n) . \end{array} \quad (18)$$

Additionally, each data property s_i has been declared as functional.

The set of assertions in \mathcal{K}_D is built in the following way. For each record r of the form $\langle v_1, \dots, v_n, T \rangle$, we create a new individual a_r and add the axioms

$$\begin{array}{l} a_r:T \\ a_r:\exists s_i. =_{v_i} \quad (i = 1..n) \end{array} \quad (19)$$

to \mathcal{K}_D . For instance, for the iris dataset described above, that has three target classes `Iris – setosa`,

Iris – versicolor and *Iris – virginica*, the KB contains the axioms

```

Iris – setosa  $\sqsubseteq$  class
Iris – versicolor  $\sqsubseteq$  class
Iris – virginica  $\sqsubseteq$  class
class  $\sqsubseteq$   $\exists$ sepal_length.double
class  $\sqsubseteq$   $\exists$ sepal_width.double
class  $\sqsubseteq$   $\exists$ sepal_length.double
class  $\sqsubseteq$   $\exists$ sepal_width.double
a1:Iris – setosa
a1: $\exists$ sepal_length. =5.1
a1: $\exists$ sepal_width. =3.5
a1: $\exists$ petal_length. =1.4
a1: $\exists$ sepal_width. =0.2 .

```

It is easily verified that the KB \mathcal{K}_D constructed for each dataset D (i) belongs to the DL $\mathcal{EL}(\mathbf{D})$ extended with functional properties; (ii) the number of classes is $|T_D| + 1$; and there are n functional datatype properties.

References

References

- [1] *Special issue on Fuzzy Clustering, Fuzzy Sets and Systems*, volume 389. Elsevier, 2020.
- [2] Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyashev. The DL-Lite family and relations. *Journal of Artificial Intelligence Research*, 36:1–69, 2009.
- [3] F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, pages 364–369, Edinburgh, UK, 2005. Morgan-Kaufmann Publishers.
- [4] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2 edition, 2007.
- [5] Liviu Badea and Shan-Hwei Nienhuys-Cheng. A refinement operator for description logics. In *Inductive Logic Programming, 10th International Conference, ILP-00*, volume 1866 of *Lecture Notes in Computer Science*, pages 40–59. Springer, 2000.
- [6] Ricardo A. Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., 1999.
- [7] James C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Springer Verlag, 1981.
- [8] Stephan Bloehdorn and York Sure. Kernel methods for mining instance data in ontologies. In *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007.*, volume 4825 of *Lecture Notes in Computer Science*, pages 58–71. Springer Verlag, 2007.
- [9] Fernando Bobillo, Marco Cerami, Francesc Esteve, Àngel García-Cerdaña, Rafael Peñaloza, and Umberto Straccia. Fuzzy description logics in the framework of mathematical fuzzy logic. In Carles Noguera Petr Cintula, Christian Fermüller, editor, *Handbook of Mathematical Fuzzy Logic, Volume 3*, volume 58 of *Studies in Logic, Mathematical Logic and Foundations*, chapter 16, pages 1105–1181. College Publications, 2015.
- [10] Fernando Bobillo and Umberto Straccia. fuzzyDL: An expressive fuzzy description logic reasoner. In *2008 International Conference on Fuzzy Systems (FUZZ-08)*, pages 923–930. IEEE Computer Society, 2008.
- [11] Fernando Bobillo and Umberto Straccia. Representing fuzzy ontologies in owl 2. In *Proceedings of the 19th IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2010)*, pages 2695–2700. IEEE Press, July 2010.
- [12] Fernando Bobillo and Umberto Straccia. Fuzzy ontology representation using OWL 2. *International Journal of Approximate Reasoning*, 52:1073–1094, 2011.

- [13] Fernando Bobillo and Umberto Straccia. The fuzzy ontology reasoner *fuzzyDL*. *Knowledge-Based Systems*, 95:12 – 34, 2016.
- [14] Fernando Bobillo and Umberto Straccia. Reasoning within fuzzy owl 2 el revisited. *Fuzzy Sets and Systems*, 351:1–40, 2018.
- [15] Stefan Borgwardt, Marco Cerami, and Rafael Peñaloza. The complexity of fuzzy \mathcal{EL} under the Lukasiewicz t-norm. *International Journal of Approximate Reasoning*, 91:179–201, 2017.
- [16] Lorenz Bühmann, Jens Lehmann, and Patrick Westphal. DI-learner - A framework for inductive learning on the semantic web. *Journal of Web Semantics*, 39:15–24, 2016.
- [17] Lorenz Bühmann, Jens Lehmann, Patrick Westphal, and Simon Bin. DI-learner structured machine learning on semantic web data. In *Companion of the The Web Conference 2018 on The Web Conference 2018, WWW 2018, Lyon , France, April 23-27, 2018*, pages 467–471. ACM, 2018.
- [18] Marco Cerami and Umberto Straccia. On the (un)decidability of fuzzy description logics under lukasiewicz t-norm. *Information Sciences*, 227:1–21, 2013.
- [19] Mahsa Chitsaz, Kewen Wang, Michael Blumenstein, and Guilin Qi. Concept learning for \mathcal{EL}^{++} ; by refinement and reinforcement. In *Proceedings of the 12th Pacific Rim international conference on Trends in Artificial Intelligence, PRICAI'12*, pages 15–26, Berlin, Heidelberg, 2012. Springer-Verlag.
- [20] Marcos E. Cintra, Maria Carolina Monard, and Heloisa de Arruda Camargo. On rule learning methods: A comparative analysis of classic and fuzzy approaches. In *Soft Computing: State of the Art Theory and Novel Applications*, volume 291, pages 89–104. Springer Verlag, 2013.
- [21] María José del Jesús, Frank Hoffmann, Luis Junco Navascués, and Luciano Sánchez. Induction of fuzzy-rule-based classifiers with evolutionary boosting algorithms. *IEEE Transactions on Fuzzy Systems*, 12(3):296–308, 2004.
- [22] Mario Drobits, Ulrich Bodenhofer, and Erich-Peter Klement. Fs-foil: an inductive learning method for extracting interpretable fuzzy descriptions. *International Journal of Approximate Reasoning*, 32(2-3):131–152, 2003.
- [23] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [24] Nicola Fanizzi. Concept induction in description logics using information-theoretic heuristics. *Int. J. Semantic Web Inf. Syst.*, 7(2):23–44, 2011.
- [25] Nicola Fanizzi and Claudia d’Amato. A declarative kernel for *ALC* concept descriptions. In *Foundations of Intelligent Systems, 16th International Symposium, ISMIS 2006, Bari, Italy, September 27-29, 2006, Proceedings*, volume 4203 of *Lecture Notes in Computer Science*, pages 322–331. Springer Verlag, 2006.
- [26] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. DL-FOIL concept learning in description logics. In Filip Zelezný and Nada Lavrač, editors, *Inductive Logic Programming*, volume 5194 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2008.
- [27] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. Induction of classifiers through non-parametric methods for approximate classification and retrieval with ontologies. *Int. J. Semantic Computing*, 2(3):403–423, 2008.
- [28] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. Induction of concepts in web ontologies through terminological decision trees. In *Machine Learning and Knowledge Discovery in Databases, European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part I*, volume 6321 of *Lecture Notes in Computer Science*, pages 442–457. Springer Verlag, 2010.
- [29] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. Towards the induction of terminological decision trees. In *Proceedings of the 2010 ACM Symposium on Applied Computing, SAC '10*, pages 1423–1427, New York, NY, USA, 2010. ACM.
- [30] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. Induction of robust classifiers for web ontologies through kernel machines. *J. Web Sem.*, 11:1–13, 2012.
- [31] Nicola Fanizzi, Giuseppe Rizzo, and Claudia d’Amato. Boosting DL concept learners. In *The Semantic Web - 16th International Conference, ESWC-19*, volume 11503 of *Lecture Notes in Computer Science*, pages 68–83, 2019.
- [32] Nicola Fanizzi, Giuseppe Rizzo, Claudia d’Amato, and Floriana Esposito. Dlfoil: Class expression learning revisited. In *Knowledge Engineering and Knowledge Management - 21st International Conference, EKAW-18*, volume 11313 of *Lecture Notes in Computer Science*, pages 98–113, 2018.
- [33] Yoav Freund and Robert E. Schapire. Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning (ICML-96)*, pages 148–156, San Francisco, CA, USA, 1996. Morgan Kaufmann Publishers Inc.

- [34] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics* 28, 28:337–374, 2000.
- [35] Ignacio Huitzil, Fernando Bobillo, Juan Gomez-Romero, and Umberto Straccia. Fudge: Fuzzy ontology building with consensuated fuzzy datatypes. *Fuzzy Sets and Systems*, 401:91–112, 2020.
- [36] Ignacio Huitzil, Umberto Straccia, Natalia Díaz-Rodríguez, and Fernando Bobillo. Datil: Learning fuzzy ontology datatypes. In *Proceedings of the 17th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 2018), Part II*, volume 854 of *Communications in Computer and Information Science*, pages 100–112. Springer, June 2018.
- [37] Josué Iglesias and Jens Lehmann. Towards integrating fuzzy logic capabilities into an ontology-based inductive logic programming framework. In *Proceedings of the 11th International Conference on Intelligent Systems Design and Applications (ISDA 2011)*, pages 1323–1328, 2011.
- [38] George J. Klir and Bo Yuan. *Fuzzy sets and fuzzy logic: theory and applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995.
- [39] S. Konstantopoulos and A. Charalambidis. Formulating description logic learning as an inductive logic programming task. In *Proceedings of the 19th IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2010)*, pages 1–7. IEEE Press, 2010.
- [40] Jens Lehmann. Hybrid learning of ontology classes. In *Machine Learning and Data Mining in Pattern Recognition, 5th International Conference, MLDM 2007, Leipzig, Germany, July 18-20, 2007, Proceedings*, volume 4571 of *Lecture Notes in Computer Science*, pages 883–898. Springer Verlag, 2007.
- [41] Jens Lehmann. DL-Learner: Learning concepts in description logics. *Journal of Machine Learning Research*, 10:2639–2642, 2009.
- [42] Jens Lehmann and Christoph Haase. Ideal downward refinement in the \mathcal{EL} description logic. In *Inductive Logic Programming, 19th International Conference, ILP 2009, Leuven, Belgium, July 02-04, 2009. Revised Papers*, volume 5989 of *Lecture Notes in Computer Science*, pages 73–87. Springer, 2009.
- [43] Jens Lehmann and Pascal Hitzler. A refinement operator based learning algorithm for the \mathcal{ALC} description logic. In *Inductive Logic Programming, 17th International Conference, ILP 2007, Corvallis, OR, USA, June 19-21, 2007, Revised Selected Papers*, volume 4894 of *Lecture Notes in Computer Science*, pages 147–160. Springer Verlag, 2007.
- [44] Jens Lehmann and Pascal Hitzler. Foundations of Refinement Operators for Description Logics. In H. Bloekel, J. Ramon, J. W. Shavlik, and P. Tadepalli, editors, *Inductive Logic Programming*, volume 4894 of *Lecture Notes in Artificial Intelligence*, pages 161–174. Springer, 2008.
- [45] Jens Lehmann and Pascal Hitzler. Concept learning in description logics using refinement operators. *Machine Learning*, 78(1-2):203–250, 2010.
- [46] Francesca A. Lisi and Donato Malerba. Ideal refinement of descriptions in al-log. In *Inductive Logic Programming: 13th International Conference, ILP-03*, volume 2835 of *Lecture Notes in Computer Science*, pages 215–232, 2003.
- [47] Francesca A. Lisi and Umberto Straccia. A logic-based computational method for the automated induction of fuzzy ontology axioms. *Fundamenta Informaticae*, 124(4):503–519, 2013.
- [48] Francesca A. Lisi and Umberto Straccia. A system for learning GCI axioms in fuzzy description logics. In *Proceedings of the 26th International Workshop on Description Logics (DL-13)*, volume 1014 of *CEUR Workshop Proceedings*, pages 760–778. CEUR-WS.org, 2013.
- [49] Francesca A. Lisi and Umberto Straccia. Can ilp deal with incomplete and vague structured knowledge? In Stephen H. Muggleton and Hiroaki Watanabe, editors, *Latest Advances in Inductive Logic Programming*, chapter 21, pages 199–206. World Scientific, 2014.
- [50] Francesca A. Lisi and Umberto Straccia. Learning in description logics with fuzzy concrete domains. *Fundamenta Informaticae*, 140(3-4):373–391, 2015.
- [51] Francesca Alessandra Lisi. Logics in machine learning and data mining: Achievements and open issues. In *Proceedings of the 34th Italian Conference on Computational Logic, Trieste, Italy, June 19-21, 2019.*, volume 2396 of *CEUR Workshop Proceedings*, pages 82–88. CEUR-WS.org, 2019.
- [52] Francesca Alessandra Lisi and Umberto Straccia. An inductive logic programming approach to learning inclusion axioms in fuzzy description logics. In *26th Italian Conference on Computational Logic (CILC-11)*, volume 810, pages 57–71. CEUR Electronic Workshop Proceedings, 2011.
- [53] Francesca Alessandra Lisi and Umberto Straccia. Towards learning fuzzy dl inclusion axioms. In *9th International Workshop on Fuzzy Logic and Applications (WILF-11)*, volume 6857 of *Lecture Notes in Computer Science*, pages 58–66, Berlin, 2011. Springer Verlag.

- [54] Francesca Alessandra Lisi and Umberto Straccia. Dealing with incompleteness and vagueness in inductive logic programming. In *28th Italian Conference on Computational Logic (CILC-13)*, volume 1068, pages 179–193. CEUR Electronic Workshop Proceedings, 2013.
- [55] Francesca Alessandra Lisi and Umberto Straccia. A foil-like method for learning under incompleteness and vagueness. In *23rd International Conference on Inductive Logic Programming*, volume 8812 of *Lecture Notes in Artificial Intelligence*, pages 123–139, Berlin, 2014. Springer Verlag. Revised Selected Papers.
- [56] Thomas Lukasiewicz and Umberto Straccia. Managing uncertainty and vagueness in description logics for the semantic web. *Journal of Web Semantics*, 6:291–308, 2008.
- [57] Pasquale Minervini, Claudia d’Amato, and Nicola Fanizzi. Learning terminological naive bayesian classifiers under different assumptions on missing knowledge. In *Proceedings of the 7th International Workshop on Uncertainty Reasoning for the Semantic Web (URSW-11)*, volume 778 of *CEUR Workshop Proceedings*, pages 63–74. CEUR-WS.org, 2011.
- [58] Pasquale Minervini, Claudia d’Amato, and Nicola Fanizzi. Learning probabilistic description logic concepts: Under different assumptions on missing knowledge. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing, SAC ’12*, pages 378–383, New York, NY, USA, 2012. ACM.
- [59] Pasquale Minervini, Claudia d’Amato, and Nicola Fanizzi. Learning terminological bayesian classifiers - A comparison of alternative approaches to dealing with unknown concept-memberships. In *Proceedings of the 9th Italian Convention on Computational Logic, Rome, Italy, June 6-7, 2012*, volume 857 of *CEUR Workshop Proceedings*, pages 191–205, 2012.
- [60] Pasquale Minervini, Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. Learning probabilistic description logic concepts under alternative assumptions on incompleteness. In *Uncertainty Reasoning for the Semantic Web III - ISWC International Workshops, URSW 2011-2013, Revised Selected Papers*, volume 8816 of *Lecture Notes in Computer Science*, pages 184–201, 2014.
- [61] Boris Motik and Riccardo Rosati. A faithful integration of description logics with logic programming. In *Proceedings of the 20th international joint conference on Artificial intelligence*, pages 477–482, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- [62] Matthias Nickles and Achim Rettinger. Interactive relational reinforcement learning of concept semantics. *Machine Learning*, 94(2):169–204, 2014.
- [63] Richard Nock and Frank Nielsen. A real generalization of discrete AdaBoost. In *17th European Conference on Artificial Intelligence (ECAI-06)*, volume 141 of *Frontiers in Artificial Intelligence and Applications*, pages 509–515. IOS Press, 2006.
- [64] Richard Nock and Frank Nielsen. A real generalization of discrete AdaBoost. *Artificial Intelligence Journal*, 171(1):25–41, 2007.
- [65] José Otero and Luciano Sánchez. Induction of descriptive fuzzy classifiers with the logitboost algorithm. *Soft Computing*, 10(9):825–835, 2006.
- [66] OWL 2 Web Ontology Language Document Overview. <https://www.w3.org/TR/owl2-overview/>. W3C, 2009.
- [67] Ana M. Palacios, Luciano Sánchez, and Inés Couso. Using the AdaBoost algorithm for extracting fuzzy rules from low quality data: Some preliminary results. In *FUZZ-IEEE 2011, IEEE International Conference on Fuzzy Systems, Taipei, Taiwan, 27-30 June, 2011, Proceedings*, pages 1263–1270. IEEE, 2011.
- [68] Luc De Raedt. *Logical and relational learning*. Cognitive Technologies. Springer, 2008.
- [69] Luc De Raedt and Kristian Kersting. Statistical relational learning. In *Encyclopedia of Machine Learning and Data Mining*, pages 1177–1187. Springer, 2017.
- [70] Achim Rettinger, Uta Lösch, Volker Tresp, Claudia d’Amato, and Nicola Fanizzi. Mining the semantic web - statistical learning for next generation knowledge bases. *Data Mining and Knowledge Discovery*, 24(3):613–662, 2012.
- [71] Giuseppe Rizzo, Claudia d’Amato, and Nicola Fanizzi. On the effectiveness of evidence-based terminological decision trees. In *Foundations of Intelligent Systems - 22nd International Symposium, ISMIS-15*, volume 9384 of *Lecture Notes in Computer Science*, pages 139–149, 2015.
- [72] Giuseppe Rizzo, Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. Tackling the class-imbalance learning problem in semantic web knowledge bases. In *Knowledge Engineering and Knowledge Management - 19th International Conference, EKAW-14*, volume 8876 of *Lecture Notes in Computer Science*, pages 453–468, 2014.

- [73] Giuseppe Rizzo, Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. Towards evidence-based terminological decision trees. In *Information Processing and Management of Uncertainty in Knowledge-Based Systems - 15th International Conference, IPMU 2014, Montpellier, France, July 15-19, 2014, Proceedings, Part I*, volume 442 of *Communications in Computer and Information Science*, pages 36–45. Springer, 2014.
- [74] Giuseppe Rizzo, Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. Inductive classification through evidence-based models and their ensembles. In *The Semantic Web. Latest Advances and New Domains - 12th European Semantic Web Conference, ESWC-15*, volume 9088 of *Lecture Notes in Computer Science*, pages 418–433, 2015.
- [75] Giuseppe Rizzo, Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. Tree-based models for inductive classification on the web of data. *Journal of Web Semantics*, 45:1–22, 2017.
- [76] Giuseppe Rizzo, Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. Approximate classification with web ontologies through evidential terminological trees and forests. *International Journal of Approximate Reasoning*, 92:340–362, 2018.
- [77] Giuseppe Rizzo, Nicola Fanizzi, Jens Lehmann, and Lorenz Böhmann. Integrating new refinement operators in terminological decision trees learning. In *Knowledge Engineering and Knowledge Management - 20th International Conference, EKAW-16*, volume 10024 of *Lecture Notes in Computer Science*, pages 511–526, 2016.
- [78] Luciano Sánchez and José Otero. Boosting fuzzy rules in classification problems under single-winner inference. *International Journal of Intelligent Systems*, 22(9):1021–1034, 2007.
- [79] Robert E. Schapire and Yoram Singer. Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3):297–336, 1999.
- [80] Mathieu Serrurier and Henri Prade. Improving expressivity of inductive logic programming by learning different kinds of fuzzy rules. *Soft Computing*, 11(5):459–466, 2007.
- [81] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [82] D. Shibata, N. Inuzuka, S. Kato, T. Matsui, and H. Itoh. An induction algorithm based on fuzzy logic programming. In N. Zhong and L. Zhou, editors, *Methodologies for Knowledge Discovery and Data Mining, Third Pacific-Asia Conference, PAKDD-99, Beijing, China, April 26-28, 1999, Proceedings*, volume 1574 of *Lecture Notes in Computer Science*, pages 268–273. Springer, 1999.
- [83] Umberto Straccia. Description logics with fuzzy concrete domains. In Fahiem Bachus and Tommi Jaakkola, editors, *21st Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 559–567, Edinburgh, Scotland, 2005. AUAI Press.
- [84] Umberto Straccia. *Foundations of Fuzzy Logic and Semantic Web Languages*. CRC Studies in Informatics Series. Chapman & Hall, 2013.
- [85] Umberto Straccia and Matteo Mucci. pFOIL-DL: Learning (fuzzy) \mathcal{EL} concept descriptions from crisp OWL data using a probabilistic ensemble estimation. In *Proceedings of the 30th Annual ACM Symposium on Applied Computing (SAC-15)*, pages 345–352, Salamanca, Spain, 2015. ACM.
- [86] Patrick Westphal, Lorenz Böhmann, Simon Bin, Hajira Jabeen, and Jens Lehmann. SML-bench - A benchmarking framework for structured machine learning. *Semantic Web*, 10(2):231–245, 2019.
- [87] Hong yang Zhu, Yi Ding, Hong Gao, and Wei Liu. Fuzzy prediction in classification of AdaBoost algorithm. In *International Conference on Oriental Thinking and Fuzzy Logic*, volume 443 of *Advances in Intelligent Systems and Computing*, pages 129–136. Springer, 2016.
- [88] L. A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.
- [89] Man Zhu, Zhiqiang Gao, Jeff Z. Pan, Yuting Zhao, Ying Xu, and Zhibin Quan. Tbox learning from incomplete data by inference in belnet+. *Knowledge-Based Systems*, 75:30–40, 2015.