

International Journal of Semantic Computing
Vol. 6, No. 4 (2012) 1–18
©World Scientific Publishing Company
DOI: 10.1142/S1793351X1200158X



NUMERIC PREDICTION ON OWL KNOWLEDGE BASES THROUGH TERMINOLOGICAL REGRESSION TREES

NICOLA FANIZZI*, CLAUDIA D'AMATO,
FLORIANA ESPOSITO and PASQUALE MINERVINI

*Dipartimento di Informatica, Università degli studi di Bari
Campus Universitario, Via Orabona 4, 70125 Bari, Italy*

**nicola.fanizzi@uniba.it*

http://lacam.diuniba.it

In the context of semantic knowledge bases, among the possible problems that may be tackled by means of data-driven inductive strategies, one can consider those that require the prediction of the unknown values of existing numeric features or the definition of new features to be derived from the data model. These problems can be cast as regression problems so that suitable solutions can be devised based on those found for multi-relational databases. In this paper, a new framework for the induction of logical regression trees is presented. Differently from the classic logical regression trees and the recent fork of the terminological classification trees, the novel *terminological regression trees* aim at predicting continuous values, while tests at the tree nodes are expressed with Description Logic concepts. They are intended for multiple uses with knowledge bases expressed in the standard ontology languages for the Semantic Web. A top-down method for growing such trees is proposed as well as algorithms for making predictions with the trees and deriving rules. The system that implements these methods is experimentally evaluated on ontologies selected from popular repositories.

Keywords: Regression tree; knowledge bases; prediction.

1. Introduction

Next generation knowledge bases will likely rely on Web-scale distributed repositories of resources. These will be indexed in terms of shared ontologies that will be expressed through standard machine-interpretable representations. Indeed, among its various facets, the *Semantic Web* is a *Web of data*. A growing number of structured data sources are distributed over the Web and comply with the new standards for data integration. Semantic Web technologies provide infrastructures for novel applications to be enabled by the possibility of querying, making inferences, etc. on such knowledge bases.

An interesting problem may concern making predictions on numeric features of the resources contained in such knowledge bases. Learning predictive models it is possible to extend the factual knowledge therein and automatically derive rules to

1 enrich the reasoning capacities, e.g. learning numeric functions (such as ranking
2 functions, reinforcement functions for planning, etc.) that are hard to express
3 analytically.

4 5 **1.1. Scenarios**

6
7 In this paper, we focus on regression problems in the context of Web ontologies. Such
8 problems naturally emerge in various scenarios. In a possible scenario, suppose a
9 (human/software) agent is looking for the distance to a certain place of interest. It
10 may query a knowledge-based system that contains a repository of spatial annota-
11 tions about places and may contain a query-answering component which may be also
12 able to perform spatial reasoning. Would the query system be able to provide an
13 answer when required information to answer the query is missing or non-derivable^a
14 through automated reasoning procedures?

15 An even more complex scenario is one where the numeric feature to be found for a
16 given resource is not analytically formalized in the knowledge base through a proper
17 relationship. Suppose an organization wants to predict the share of a given TV
18 program given a categorization of TV programs and a history of shares registered
19 in the past for other programs (see also the next Ex. 1, where this scenario is
20 formalized). Scoring a certain share can be hardly modeled with a relationship in the
21 knowledge base, i.e. its definition can be only extensional: an (incomplete) enumer-
22 ation of cases observed in the past.

23 Determining or predicting the filler of a real-valued role is not a standard inference
24 service for DL reasoners. Even quantifying the degree of membership with respect to
25 a given query concept may be cast as a regression problem (especially when reasoner
26 is not able to provide a positive or negative answer). An inductive inference service
27 that is able to at least suggest an answer to these problems may enhance the query
28 answering components of semantic repository managers with further reasoning
29 capabilities.

30 31 **1.2. Solutions**

32
33 The induction of decision trees is among the most well-known machine learning
34 techniques, along with its extensions towards logical representations in clausal form
35 [14, 16, 3]. Together with several adaptations of classical learning algorithms based
36 on refinement operators for inducing DL concepts (e.g. see [17]), the general tree
37 induction framework has recently been extended to cope with the DL languages
38 supporting the Web ontologies [7]: the tests at the internal nodes of a logical
39 decision tree are represented through DL concepts (class expressions). Compared to
40 the other logic representations of the tree tests, ours can naturally comply with the
41 semantics of the data expressed in the standard languages for the Semantic Web

42
43 ^aSpatial reasoning systems are generally not able to perform inductive inferences.

1 context. The resulting *terminological decision trees* have been used as alternative
2 models for classifying individuals and/or for inducing new concept descriptions from
3 examples.

4 Following [16], the setting is further extended to address regression rather than
5 classification problems: leaf-nodes have to indicate the values for a target function to
6 be learned. These values may be provided as constants (real numbers) or even
7 through more complex parametric models (e.g. linear functions) to be applied to the
8 individuals routed to the given leaf.

9 We introduce a new type of logical model trees called *Terminological Regression*
10 *Trees*. We propose a tree-induction algorithm that adopts a classical top-down
11 *divide-and-conquer* strategy with the use of refinement operators for DL concept
12 descriptions [13, 17, 6]. An *ad hoc* evaluation measure is adopted as a heuristic for
13 deciding the test concepts that are installed at the internal nodes.

14 In an extensive experiment, the implementation of the proposed algorithms was
15 applied to real ontologies selected from popular repositories. Artificial problems were
16 crafted by considering the prediction of the values of a continuous target function
17 through trained regression trees. The experiments show empirically the feasibility of
18 the method as very limited errors were observed on average.

19 First-order trees may be considered for various applications such as clustering,
20 categorization or numeric prediction. In particular, the induction of regression trees
21 may be considered an alternative way for learning *ranking functions*, a problem that
22 has also been considered in the context of DL knowledge bases [8].

23 24 **1.3. Plan**

25
26 In the remainder of this paper, we briefly introduce the basics of the underlying
27 representation (Sec. 2). Then, the learning problem is formalized (Sec. 3) and we
28 present a solution based on terminological regression trees presenting the algorithms
29 for growing them, deriving rules and for making predictions (Sec. 4). Experiments
30 proving the effectiveness of the approach are reported in Sec. 5. Finally, possible
31 applications and further developments are discussed in Sec. 6.

32 33 **2. DL Knowledge Bases**

34
35 We are targeting OWL knowledge bases with no commitment to a specific version or
36 dialect. In order to make the paper self-contained, we shortly recall the essentials of
37 the representation and reasoning adopted. More details can be easily found by
38 consulting the reference handbook for DLs [1].

39 Roughly, in the terminological formalisms *concepts* and *relations* are used to
40 describe classes of resources in a domain and relationships between them. Primitive
41 *concepts* $N_C = \{C, D, \dots\}$ are interpreted as subsets of a domain of objects
42 (resources) and primitive *roles* $N_R = \{R, S, \dots\}$ are interpreted as binary relations
43

1 on such a domain (properties). *Individuals* represent the objects through names
 2 chosen from the set $N_I = \{a, b, \dots\}$. The meaning of the descriptions is defined by an
 3 *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is the *domain* of the interpretation and the
 4 functor $\cdot^{\mathcal{I}}$ stands for the *interpretation function*, mapping each individual a to some
 5 $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ and hence, the *intension* of concepts and roles to their *extension* (respec-
 6 tively, a subset of $\Delta^{\mathcal{I}}$ and a binary relation defined on $\Delta^{\mathcal{I}}$).

7 Complex concept descriptions are built using atomic concepts (including the *top*
 8 and *bottom concepts* denoted, resp., with \top and \perp) and primitive roles by means of
 9 specific constructors. In \mathcal{ALC} the following constructors are allowed: *full concept*
 10 *negation* ($\neg C$), concept *conjunction* (denoted with $C \sqcap C'$) and, then also *disjunc-*
 11 *tion* (denoted with $C \sqcup C'$), and the *existential restriction* and the *value restriction*
 12 on roles, denoted, resp. with $\exists R.C$ and $\forall R.C$, selecting the individuals in the domain
 13 related through R to (some, resp. all) individuals that belong to C .

14 Additional constructors extend the expressiveness of the \mathcal{ALC} language. Besides,
 15 concrete domains (\mathbf{D}) with their specific semantics can be dealt with. They may
 16 include basic data types, such as numerical types, strings, etc., but also more complex
 17 types, such as tuples of the relational calculus or time intervals. In this paper, the
 18 interest is limited to the case of real values, in case one wants to encode the learning
 19 problem in the same language.

20 The notion of *subsumption* between concepts is given in terms of the interpret-
 21 ations: Given two concept descriptions C and D , C is *subsumed* by D , denoted by
 22 $C \sqsubseteq D$, iff for every interpretation \mathcal{I} of \mathcal{T} it holds that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. Hence, $C \equiv D$
 23 amounts to $C \sqsubseteq D$ and $D \sqsubseteq C$. The interpretations of interest will be limited to
 24 those satisfying the axioms in the knowledge base. A *knowledge base* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$
 25 contains two components: a TBox \mathcal{T} and an ABox \mathcal{A} . \mathcal{T} is a set of terminological
 26 axioms $C \sqsubseteq D$, yet we will consider only inclusions $A \sqsubseteq D$, where $A \in N_C$ is a con-
 27 cept name (atomic) and D is a concept description given in terms of the language
 28 constructors. The ABox \mathcal{A} contains extensional assertions (ground facts) on concepts
 29 and roles, e.g. $C(a)$ and $R(a, b)$, meaning, respectively, that $a^{\mathcal{I}} \in C^{\mathcal{I}}$ and
 30 $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. An interpretation satisfying all the axioms in the knowledge base is
 31 said to be a *model* for it. Hence, the usual notions of satisfiability, consistency, etc.
 32 apply also for these logics.

33 The most important inference service from the inductive point of view is *instance*
 34 *checking* [1] which amounts to ascertain class-membership assertions: $\mathcal{K} \models C(a)$,
 35 where \mathcal{K} is the knowledge base a is an individual name and C is a concept definition
 36 given in terms of the concepts accounted for in \mathcal{K} .

37 An important difference with clausal logic (multi relational databases), where the
 38 logic decision trees have stemmed from, is the *open-world assumption* (OWA) which
 39 has consequences on answering to class-membership queries. Thus it may happen
 40 that an object that cannot be proved to belong to a certain concept is not necessarily
 41 a counterexample for that concept. That would only be interpreted as a case of
 42 insufficient (incomplete) knowledge for proving the assertion.

3. Regression Problems with DL Knowledge Bases

Given a DL knowledge base as the source for the background knowledge and a target real-valued function (which may have been encoded as a functional role) whose analytic form is unknown (or too complex to be expressed), one may suppose that domain experts are able to provide the values of such a function for a limited number of individuals, e.g. in the form of role assertions.

In this setting, the objective is then to induce an (approximated) analytic function which can exhibit the same behavior of the target function on the training individuals and predict approximately correct values for further ones. More formally:

Definition 1. Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a DL knowledge base, with $\text{Ind}(\mathcal{A})$ denoting the individuals occurring in \mathcal{A} .

Given

- a target function $f : \text{Ind}(\mathcal{A}) \rightarrow \mathbf{R}$
(or a functional role R ranging on the concrete domain \mathbf{R})
whose analytic form is unknown;
- a sample of (training) individuals for which the f -value is known, i.e. f (resp. R) may be partially (extensionally) defined:
 $\mathcal{S}(\mathcal{A}) = \{(a, f(a)) \mid a \in \text{Ind}(\mathcal{A})\} \subseteq \text{Ind}(\mathcal{A})$
(resp. $\mathcal{S}(\mathcal{A}) = \{(a, v) \mid R(a, v) \in \mathcal{A}\}$)
- a small $\varepsilon > 0$

Build a regression model $h : \text{Ind}(\mathcal{A}) \rightarrow \mathbf{R}$ so that:

$$|h(a) - f(a)| < \varepsilon, \quad \forall a \in \text{Ind}(\mathcal{A})$$

One may consider the case of relevance feedback where the evaluation of Web resources may be exploited to learn an analytical model which adapts as new resources become available. Preference learning in the context of ontologies is another closely related task.

An example on the domain of *TV Programs* may better illustrate the learning task as follows:

Example 1. (TV programs) Suppose a knowledge base concerning TV Programs is available whose terminology allows for expressing information about the programs, such as their type, production year, broadcast date, etc. Suppose one wants to predict the share, i.e. to learn a share function based on the available terminology and assertions known about the TV programs. The TBox \mathcal{T} may include the background knowledge shown in Fig. 1.

Besides, since some concepts are meant to be disjoint, suitable axioms must be added to \mathcal{T} :

$$\begin{aligned} \{\text{Film} \sqsubseteq \neg(\text{TalkShow} \sqcup \text{Serial}), \\ \text{Show} \sqsubseteq \neg(\text{Film} \sqcup \text{Serial}), \\ \text{SoapOpera} \sqsubseteq \neg(\text{SitCom} \sqcup \text{Film}), \dots\} \end{aligned}$$

6 *N. Fanizzi et al.*

```

1      {Program  $\sqsubseteq \forall$ on.TVNetwork  $\sqcap \forall$ producedIn.Year  $\sqcap \forall$ broadcastOn.Date,
2      News  $\sqsubseteq$  Program  $\sqcap \forall$ hasAnchor.Person,
3      Entertainment  $\sqsubseteq$  Program,
4      Cultural  $\sqsubseteq$  Program,
5      Film  $\sqsubseteq$  (Movie  $\sqcup$  Documentary)  $\sqcap \forall$ directedBy.Person,
6      Movie  $\sqsubseteq$  (Entertainment  $\sqcup$  Cultural)  $\sqcap \forall$ starring.Actor,
7      Documentary  $\sqsubseteq$  Film  $\sqcap$  Cultural  $\sqcap \forall$ starring.⊥,
8      Serial  $\sqsubseteq$  Entertainment  $\sqcap \neg$ News  $\sqcap \forall$ starring.Actor,
9      SoapOpera  $\sqsubseteq$  Serial,
10     SitCom  $\sqsubseteq$  Serial,
11     Show  $\sqsubseteq$  Entertainment  $\sqcap \forall$ hostedBy.Person,
12     TalkShow  $\sqsubseteq$  Show  $\sqcap$  News  $\sqcap \forall$ dealsWith.Topic}  $\subseteq \mathcal{T}$ 

```

Fig. 1. Background knowledge for the TV-rating example.

The ABox may contain the following assertions:

```

16     { SitCom(SEINFELD), on(SEINFELD, NBC),
17     broadcastOn(SEINFELD, 20101009),
18     starring(SEINFELD, JL_DREYFUS),
19     TalkShow(LATESHOW), on(LATESHOW, CBS),
20     hostedBy(LATESHOW, D_LETTERMAN),
21     broadcastOn(LATESHOW, 20100212),
22     Documentary(SICKO), on(SICKO, CURRENT),
23     directedBy(SICKO, M_MOORE),
24     broadcastOn(SICKO, 20101111), ... }  $\subseteq \mathcal{A}$ 

```

Suppose the intended function to be predicted is the share of a given program. A sample of this function may contain the following couples-

```

29      $\mathcal{S}(\mathcal{A}) = \{ (\text{SEINFELD}, 21.5\%), (\text{LATESHOW}, 11.2\%),$ 
30      $(\text{SICKO}, 23.1\%), \dots \}$ 

```

Note that this problem differs from settings where the aim is building a classifier through logical or statistical methods employing, for example, *support vector machines* [4], which may be tackled through more suitable kernel machines (for regression). Further related settings will not be discussed further.

4. Growing and Exploiting Terminological Regression Trees

In the context of the clausal representations adopted in *inductive logic programming* (ILP), *first-order logical trees* (FOLTs) are defined [3] as binary trees in which

- (1) the inner nodes contain tests in the form of conjunctions of literals;
- (2) left and right branches stand, resp., for the truth-value (resp. true and false) determined by the test evaluation;

(3) different nodes may share variables, yet a variable that is introduced in a certain node must not occur in the right branch of that node.

In the case of *regression trees* [14–16], the leaf-nodes contain a local regression model in one of the forms that will be discussed later.

We extend the original definition along various directions:

Definition 1. *Terminological Regression Trees* (TRTs) are FOLTs whose internal nodes contain DL concept descriptions as tests and whose leaf-nodes contain functions (*local models*) to be evaluated to get a predicted value.

In the simplest form leaf-nodes represent constant functions that return average values for all the instances that reach those leaves during the evaluation phase (see below). Figure 2 shows a TRT in the context of the knowledge base proposed in Example 1.

We will consider a generic binary tree structure whose nodes N are tuples containing the following fields:

- test concept description: $N.test$;
- set of training instances rooted at the node: $N.set$;
- regression model: $N.model$;
- references to the subtrees, resp. $N.left$ and $N.right$.

We now show how to grow TRTs and how to use them for predicting values.

4.1. Induction of TRTs

While DL concept induction algorithms generally adopt a separate-and-conquer covering strategy [13, 6], the tree learning procedures adopt a divide-and-conquer strategy.

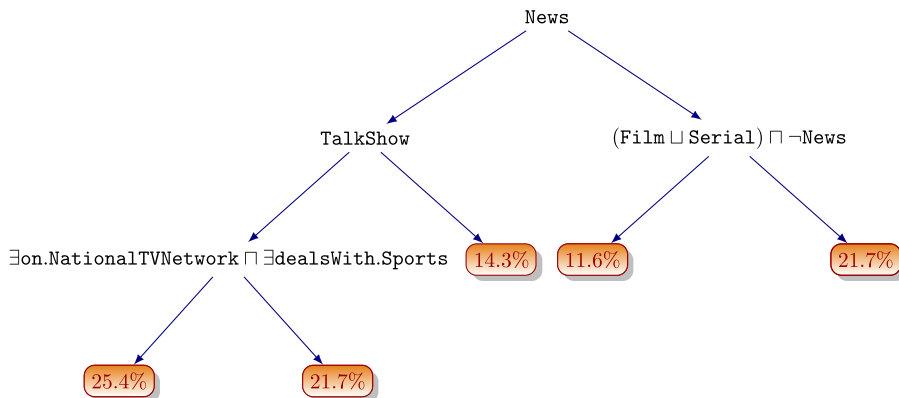


Fig. 2. A simple example of a TRT referring to the problem introduced in Example 1.

1 In the case of TDTs [7], the FOLT induction procedure is adapted to deal with
 2 internal nodes that contain DL concepts to be interpreted w.r.t. their semantics.
 3 Since the choice of such nodes involves a specialization task this may depend on the
 4 properties of the particular space of concept descriptions. The other difference is
 5 in the policy for labeling the leaf-nodes. While for decision trees it suffices to indicate
 6 the target class for the instances that are routed to those nodes, for regression trees
 7 the (regression) model needed for the local regression problem should be decided [18].

8 As regards the first issue, the subsumption relationship \sqsubseteq induces a partial order
 9 on the space of DL concept descriptions. Then, a specialization task can be cast as a
 10 search in such a partially ordered space. In such a setting, suitable operators to
 11 traverse the search space are required [17]. As this space is very large (especially from
 12 a syntactic viewpoint) and may turn out to exhibit a lot of redundancy (equivalent
 13 concepts), also depending on the expressiveness of the underlying DL logic, some bias
 14 is needed to constrain the search. In [7] only refinement operators working on \mathcal{ALC}
 15 constructors are used, even in the context of ontologies represented through more
 16 expressive languages, thus trading completeness for efficiency. A further constraint
 17 that can be considered leads to prefer candidate specializations for which the training
 18 instances tend to exhibit a definite membership, discouraging those concepts for
 19 which individuals show an unknown membership (due to the OWA).

20 The TRT-induction procedure adapts the schema implemented by the extensions
 21 of TILDE towards first order regression [3]. A sketch of the main procedure is reported
 22 as Algorithm 1. It reflects the standard tree induction algorithms with the addition of
 23

24 **Algorithm 1** The main growing procedure for terminological regression trees
 25 $\text{GROWTRT}(\mathcal{K}, C, LS): T$;

26 **Input:**

27 \mathcal{K} : knowledge base,
 28 C : concept description,
 29 LS : local set of training individuals

30 **Output:** T : TRT

31 1: Initialize new tree T
 32 2: $N.\text{set} \leftarrow LS$
 33 3: **if** $|LS| < m$ **or** $\text{sd}(LS)/TSD \leq \theta$ **then**
 34 4: $N.\text{model} \leftarrow \text{BUILDREGMODEL}(LS)$
 35 5: $N.\text{left} \leftarrow \text{nil}; N.\text{right} \leftarrow \text{nil}$
 36 6: **else**
 37 7: $SPC \leftarrow \text{GENERATESPECIALIZATIONS}(\mathcal{K}, C, LS)$
 38 8: $C^* \leftarrow \text{SELECTBESTSDR}(\mathcal{K}, SPC, LS)$
 39 9: $(LS_+, LS_-) \leftarrow \text{SPLIT}(C^*, LS)$
 40 10: $N.\text{test} \leftarrow C^*$
 41 11: $N.\text{left} \leftarrow \text{GROWTRT}(\mathcal{K}, \text{SIMPLIFY}(C \sqcap C^*), LS_+)$
 42 12: $N.\text{right} \leftarrow \text{GROWTRT}(\mathcal{K}, \text{SIMPLIFY}(C \sqcap \neg C^*), LS_-)$
 43 13: **end if**
 44 14: $T.\text{root} \leftarrow N$
 45 15: **return** T

1 the treatment of unlabeled training individuals. The procedure can be invoked
 2 passing \top as starting concept and a set of individuals containing all training
 3 examples (for which the value of the target function is known).

4 The initial conditional statement takes care of the base case for the recursion,
 5 namely when a limited number of individuals got sorted to the current subtree root-
 6 node (w.r.t. some parameter m) or the local standard deviation ($\text{sd}(LS)$) is low w.r.t
 7 the global one (TSD), i.e. their ratio is less than a threshold θ , then the resulting leaf
 8 value is decided on the grounds of the preferred local regression algorithm, installing
 9 the resulting regression model $h(\cdot)$ based on the values in the local set of individuals
 10 $LS = \{a_1, \dots, a_n\}$ [18]:

- 11
 12 (1) a constant function, i.e. a real value v that averages the values of the training
 13 instances that placed at the leaf node: $h(x) = v = \sum_{i=1}^n f(a_i)/|LS|$;
 14 (2) the value that may be determined on-the-fly, for a given instance x , by means
 15 of a simple instance-based regression procedure:

$$16 \quad h(x) = \frac{\sum_{i=1}^n w_i f(a_i)}{\sum_{i=1}^n w_i} \quad (1)$$

17
 18 where the weights may be selected as inversely proportional to the distance,^b e.g.
 19 $w_i = [d(x, a_i)]^{-2}$;

- 20
 21 (3) extending the idea of the previous point, a regression algorithm that can operate
 22 with a limited number of training instances (e.g. a *radial basis function network*),
 23 may approximate the value computation as follows:
 24

$$25 \quad h(x) = \frac{\sum_{i=1}^n f(a_i) \kappa(d(x, a_i))}{\sum_{i=1}^n \kappa(d(x, a_i))} \quad (2)$$

26
 27 using a Gaussian kernel κ , for instance.
 28
 29

30 In the following recursive part of the algorithm (lines 7–12) a list *SPC* of (satisfiable)
 31 candidate concept description are randomly generated (line 7), that can specialize
 32 the current description C . As mentioned above, the generation is constrained to
 33 produce concepts that can really discriminate the individuals in LS , so that they can
 34 contribute as good node tests.

35 Then (line 8), the best one (C^*) is selected (invoking `SELECTBESTSDR`) in terms of
 36 a purity criterion based on the subsets of individuals resulting from the membership/
 37 non-membership determined by the given test concept description. The measure is
 38 derived from *information gain*. In the DL setting the problem is made more complex
 39 by the presence of instances which cannot be labeled as members or non-members
 40 w.r.t. the given concept. Their contribution may be considered as proportional to the
 41 prior distribution of positive and negative examples. We propose the following
 42

43 ^bSee similarity or distance functions surveyed in [5], for example.

10 *N. Fanizzi et al.*

1 *standard deviation reduction* measure:

$$2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \quad 8 \quad 9 \quad \text{SDR}(C, S) = c_0 \left(\text{std_dev}(S) - \sum_{S' \in \{S_+, S_-\}} \frac{|S'|}{|S|} \text{std_dev}(S') \right) \quad (3)$$

6 where S is divided into S_+ , S_- and S_0 depending on the instance check of the individuals w.r.t. C (resp. membership, non-membership, uncertain membership subsets) and c_0 is a discount factor for the missing values, $c_0 = |S \setminus S_0|/|S|$.

9 Once the best description C^* has been selected, it is installed (line 10) as the current subtree root and the sets of individuals sorted to this node are subdivided (line 9) according to their classification w.r.t. such a concept. Note that unlabeled individuals must be sorted to both subtrees.

11 Finally the construction of the left and right subtrees is requested recursively (lines 11–12), passing the sets of individuals resulting from the split and the concept descriptions $C \sqcap C^*$ and $C \sqcap \neg C^*$ in their simplified version. Function SIMPLIFY is meant to apply normalization and simplification rules reducing the structural complexity of the concepts.

19 4.2. Prediction

20 The TRTs are to be used for predicting the values of numeric functions for given individuals. Algorithm 2 illustrates the related procedure that is supported by some auxiliary functions: LEAF() to determine whether a node is a leaf of the argument tree, ROOT() which returns the root node of the tree passed as argument, and EVAL() which returns evaluation w.r.t. the given model.

27 **Algorithm 2** Numeric prediction with TRTs:

28 PREDICT(a, T, \mathcal{K}): \mathbf{R}

29 **Input:**

30 a : individual,
31 T : TRT,
32 \mathcal{K} : knowledge base

32 **Output:** $v : \mathbf{R}$

33 1: $N \leftarrow \text{ROOT}(T)$;
34 2: **while not** LEAF(N) **do**
35 3: $C \leftarrow N.\text{test}$;
36 4: **if** $\mathcal{K} \models C(a)$ **then** {*follow left-branch*}
37 5: $N \leftarrow \text{ROOT}(N.\text{left})$
38 6: **else if** $\mathcal{K} \models \neg C(a)$ **then** {*follow right-branch*}
39 7: $N \leftarrow \text{ROOT}(N.\text{right})$
40 8: **else** {*special case*}
41 9: **return** EVAL(default_model, $N.\text{set}$)
42 10: **end if**
43 11: **end while**
44 12: **return** EVAL($N.\text{model}$, $N.\text{set}$)

consequent depends on the type of model adopted: it may be based on the local set of training individuals and then it can be evaluated on-the-fly or it may be a parametric function whose coefficients vary from leaf-node to leaf-node.

As an example, looking back at the TRT depicted in Fig. 2, a rule definition that would be extracted for the leftmost path is: `News` \sqcap `TalkShow` \sqcap \exists on.`NationalTVNetwork` \sqcap \exists dealsWith.Sports \rightarrow 25.4%.

A rule language may be the natural candidate for encoding such list of rules in a DL program, so that a suitable reasoner may be employed for the prediction task instead of an *ad hoc* one dealing with TRTs.

5. Experimental Evaluation

Preliminary experiments with the implementation of the algorithms presented showed encouraging results. However specific standard datasets to test them are still hard to find (i.e. ontologies that are rich with numeric datatype assertions). In order to provide more than a mere proof of concept, artificial prediction problems on real ontologies to be solved by inducing TRTs have been automatically generated.

5.1. Setup

A number of OWL ontologies on various domains have been selected^c: MDM0.73, SURFACE WATER MODEL (SWM), WINE, UNITS, IMDB, FAMILY TREE (FAMILY) TRANSPORTATION (TRANS), NEW TESTAMENT NAMES (NTN), the *BioPax Glycolysis* ontology (BioPaxG), the FINANCIAL ontology and one large ontology generated by the *Lehigh University BenchMark* (LUBM). Table 1 summarizes important details concerning these ontologies, in terms of the numbers of concepts (classes), object and datatype properties and individuals.

For each ontology, artificial learning problems were created on-the-fly as follows: 20 class expressions were randomly generated by composition of 2 through 8 classes using^d the \mathcal{ALC} constructors in OWL2. For each class expression C , the target function f_C to be approximated assesses the likelihood of membership of the given individual to C . The value of $f_C : \text{Ind}(\mathcal{A}) \rightarrow [0, 1]$ for each individual x was assigned by recurring to a simple procedure based on density estimation that returned the likelihood measure:

$$\Pr(h_C(x) = +1 \mid x) = \frac{\pi_{+1}^C \hat{D}_{+1}(x)}{\pi_{-1}^C \hat{D}_{-1}(x) + \pi_{+1}^C \hat{D}_{+1}(x)} \quad (4)$$

^cThey are publicly available in well-known ontology repositories, the Protégé library (<http://protege.stanford.edu/plugins/owl/owl-library>) and the TONES repository (<http://owl.cs.manchester.ac.uk/repository>).

^dNote that since the ontologies employed in the experiment are expressed in various DL languages, see Table 1, these class expressions are represented in a more expressive language than \mathcal{ALC} .

Table 1. Facts concerning the ontologies employed in the experiments.

Ontology	DL language	#concepts	#object prop's	#datatype prop's	#individuals
UNITS	$\mathcal{ALUOF}(\mathbf{D})$	12	3	5	103
MDM0.73	$\mathcal{ALCHOF}(\mathbf{D})$	196	22	3	112
SWM	\mathcal{ALCOF}	19	9	0	115
WINE	$\mathcal{ALCOF}(\mathbf{D})$	75	12	1	161
TRANS	$\mathcal{ALCH}(\mathbf{D})$	445	89	4	183
IMDB	$\mathcal{ALIN}(\mathbf{D})$	7	5	13	302
BioPAXG	$\mathcal{ALCIF}(\mathbf{D})$	74	70	40	323
FAMILY	$\mathcal{SHIOF}(\mathbf{D})$	23	56	6	436
LUBM	$\mathcal{ALR}_+HI(\mathbf{D})$	55	36	11	555
NTN	$\mathcal{SHIF}(\mathbf{D})$	47	27	8	676
FINANCIAL	\mathcal{ALCIF}	60	16	0	1000

where h_C is a binary classification function (induced through an auxiliary method, the k -nearest neighbors [11], with k set to \sqrt{N} , N being the number of training instances for the problem), the π_v^C 's stand for prior probabilities of getting the respective classification value in $\{-1, +1\}$ and the \hat{D}_v 's are estimates of the density function for the given classification value, around the input individual x . Hence every individual in $\text{Ind}(\mathcal{A})$ was assigned a value $f_C + \varepsilon$, for each C , where $\varepsilon < .001$ is a small random perturbation. The prior distribution of positive and negative instances were computed for each ontology. A collection of couples $\langle x_i, f_C(x_i) \rangle$ constituted the dataset for a learning problem.^e

In the experiment design a 10-fold cross validation strategy was adopted. The performance was evaluated measuring the average error the value predicted for the test individuals w.r.t. the various random class expressions using the induced trees $h_C(x)$ compared to $f_C(x)$. The default settings (threshold $\theta = .05$ and $m = 5$) were considered (see Algorithm 1). The PELLET reasoner (v. 2.2.2) was employed for the required reasoning services.

5.2. Results

Due to space limitation, we can only report aggregate results of the learning problems. Table 2 shows, for each ontology, the outcomes in terms of the average error and the related standard deviation over the 10 folds averaged also over the various datasets for the class expressions generated.

Preliminarily, we found that the search procedure was quite accurate: it made few critical mistakes, especially when the concepts considered are known to have many instances in the ontology. Indeed, the overall average error is 1.39E-2 in an interval of [1.00E-6, 2.99E-2] whose extrema correspond, resp., to the values observed for the ontologies TRANS and LUBM. The standard deviations are also very limited (overall average 1.10E-2) suggesting the method was quite stable over the various learning

^eA snippet of the code for random concept generation and some of the generated datasets will be available at: <http://lacam.di.uniba.it/nico/research/ontologymining.html>.

Table 2. Results of the experiments on TRTs: average errors \pm standard deviations.

Ontology	Average error	Std. dev.
UNITS	2.22E-2	2.21E-2
MDM0.73	1.07E-2	9.93E-3
SWM	1.30E-2	1.00E-2
IMDB	1.33E-2	1.31E-2
WINE	1.01E-2	4.97E-3
TRANS	1.00E-6	1.00E-9
BioPaxG	1.89E-3	2.83E-3
LUBM	2.99E-2	2.19E-2
FAMILY	2.37E-2	2.28E-2
NTN	1.62E-2	7.86E-3
FINANCIAL	1.10E-2	5.66E-3

problems and ontologies. The cases of ontologies for which this measure was higher are probably due to the limited number of available examples.

The elapsed time was quite short (considered the current prototypical stage of the implementation, which lacks possible optimizations): for all the replications of the experiment (training+test) on the entire set of generated class expressions, they range from a few minutes to about 1.5 hours for a whole experiment on a medium sized ontology (in terms of number of individuals) including the time consumed by the reasoner for the deductive instance checks.

From a qualitative viewpoint, it must be noticed that the class expressions that may derived from the tree branches generally tend to be more complex than the original ones that were generated for the related learning problems.

The height of the resulting TRTs varies a lot among the various learning problems. This size can be controlled by tweaking the two parameters m and θ . Further experiments (not reported here) showed that lower values have the effect of increasing the height of the TRTs and, as expected, their precision (lower average error). However, since the training phase is generally more computationally expensive than predicting with the obtained TRTs a tradeoff may be made setting higher values for these parameters. This may be further extended to work in an interactive mode letting a user decide on whether to further expand a node or transform it to a leaf.

In an ontology population perspective, the predicted values are interesting because they suggest new assertions which cannot be logically derived by using a deductive reasoner yet they might be used to complete a knowledge base, e.g. after being validated by ontology engineers and domain experts.

5.3. Rank prediction in the linked user feedback LOD dataset

For verifying the applicability of the proposed TRT induction method in a real world context, it was evaluated on a real knowledge base, namely the *Linked User Feedback*

1 (LUF) dataset.^f LUF is part of an effort to semantically publishing and retrieving
2 user-generated feedbacks (such as ratings, comments and tags); as part of this effort,
3 ratings from the *Linked Movie Data Base*^g (LinkedMDB) were processed and inte-
4 grated into the CKAN^h dataset, which is part of the *Linked Open Data* (LOD) cloud.ⁱ

5 To evaluate the effectiveness and feasibility of our approach, it was applied to a
6 film ranking prediction task: given a sample of ratings provided by users, the system
7 induces a ranking rule to predict ratings for unranked movies. In order to leverage
8 the large amount of structured knowledge available through the LOD cloud, we
9 extracted a fragment of the DBpedia [2] knowledge base related to movies ranked in
10 the LUF dataset.

11 For this task, we followed the procedure described in [12]: starting from resources
12 representing movies, a search was performed in the RDF graph (with recursion depth
13 1), and up to 1000 superclasses were extracted for each reached object. Such an
14 extraction process resulted in an OWL 2 EL fragment containing 4789 concepts, 59
15 object properties and 3082 individuals.

16 To fit to the RDF(S)'s lack of negation and disjunction [9], we used a “LOD-
17 friendly” variant of TRT — during the TRT growth and prediction processes we
18 assumed that, given a splitting node associated to a concept C and a generic indi-
19 vidual a in \mathcal{K} , $\mathcal{K} \not\models C(a)$ implies that $\mathcal{K} \models \neg C(a)$ (thus assuming a consistent com-
20 pletion of the knowledge base). An effect of this approach is that e.g. during the
21 PREDICT procedure, given a generic individual a and a splitting node associated to a
22 concept C , we follow the left-branch if $\mathcal{K} \models C(a)$ and the right-branch otherwise,
23 and evaluate only at leaf nodes. At each step of the TRT growth process, the set of
24 candidate concepts to be associated to a split node was the set of all atomic concepts
25 in the ontology.

26 Figure 3 shows an example of an induced TRT, modelling the rankings given by
27 an user to movies in the LUF dataset.^j From preliminary empirical evaluations on the
28 LUF dataset we observed that, in induced TRTs, the concepts near to the root
29 tended to be meaningful (e.g. $y : \text{English-languageFilms}$, $y : \text{AmericanFilms}$,
30 $y : \text{AmericanBiographicalFilms}$, $y : \text{1990sRomanticComedyFilms}$); while, increas-
31 ing the distance from the root, there was an higher tendency in adding nodes
32 splitting apart very few individuals with a divergent ranking from the others
33 in that node, and the corresponding concepts were hardly meaningful (e.g.
34 $y : \text{FilmsBasedOnDarkHorseComics}$, $y : \text{FilmsBasedOnPlays}$). A possible cause for
35 this phenomenon is the *curse of dimensionality*: some concepts could improve the
36 SDR only because of statistical fluctuations within the data, instead of statistical
37 regularities. In future works, this could justify the use of feature selection as a
38

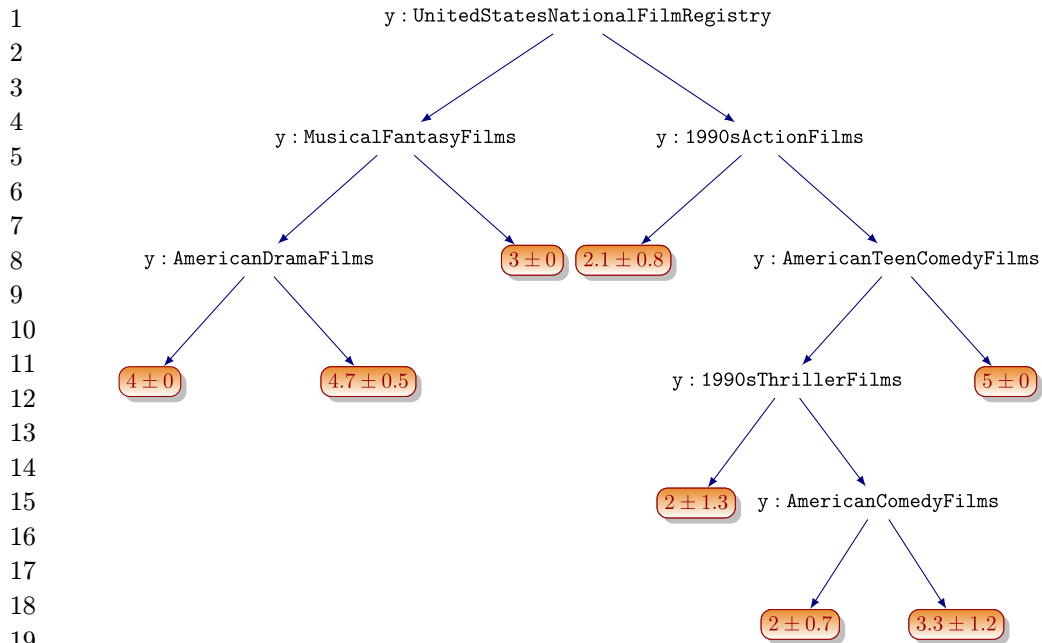
39 ^f<http://thedatahub.org/dataset/linked-user-feedback>

40 ^g<http://www.linkedmdb.org/>

41 ^h<http://ckan.org/>

42 ⁱ<http://lamboratory.com/2011/12/17/linked-user-feedback/>

43 ^j<http://soa4all.isoco.net/luf/users/movies303>. For brevity, we will use y as a prefix corresponding to the
<http://dbpedia.org/classes/yago/> namespace

16 *N. Fanizzi et al.*

23
24
25
26
27
28
29

Fig. 3. Portion of an induced TRT modelling user preferences within the Linked User Feedback (LUF) dataset.

pre-processing step prior to TRT induction (e.g. a restriction of the space of possible candidate concepts for splitting nodes); the application of growth stopping criteria and pruning procedures; and the use of scoring function different from the simple SDR [11, 10].

6. Conclusions and Developments

30
31
32
33
34
35
36
37
38
39
40

Methods for predicting numerical values in the context of DL knowledge bases were investigated. Introducing the notion of terminological regression trees, which stem from first order trees, we proposed induction algorithms as an adaptation of standard top-down growing methods complying with the issues related to the different representation. We have shown how to exploit such models to predict numerical values, such as those belonging to the range of a datatype or, generally, a complex function that is hard to define analytically. Moreover, the regression trees may be converted into rules which may be easily expressed in Semantic Web rule languages. The experiments made on various ontologies showed that the method is quite effective, its performance depending on the number (and distribution) of the available training individuals.

41
42
43

We plan to extend this work in various directions. At this stage, the adopted refinement operators search the potential candidate refinements incompletely. The methods can manage ontologies represented in more expressive languages than *ACC*

1 but reuse the concepts therein as atoms and building new ones exclusively thorough
2 *ALC* concept constructors. The problem of the dimensions of the resulting trees
3 suggest the adoption of pruning strategies after the induction phase or interleaving
4 with it which could result in more compact and better predicting trees, even in the
5 presence of large number of training instances. Better indices for the standard devi-
6 ations shall be explored, especially to better take into account the concept simplicity
7 or the uncertainty related to the unlabeled individuals (and missing values).

8 Possible applications may include ranking queried resources and the enrichment of
9 an ontology with inductively annotated assertions for roles ranging on concrete
10 domains (numeric datatypes). Finally, the presented trees may be the basis for
11 alternative hierarchical clustering algorithms where clusters would be formed
12 grouping individuals on the grounds of the invented subconcept instead of their
13 similarity, as this may be hardly defined with such complex representations.

14 References

- 15
- 16
- 17 [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi and P. Patel-Schneider (eds.), *The*
18 *Description Logic Handbook* (Cambridge University Press, 2003).
- 19 [2] C. Bizer, J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak and S. Hellmann,
20 *DBpedia — a crystallization point for the web of data*, *Web Semant.* **7**(3) (2009)
21 154–165.
- 22 [3] H. Blockeel, L. De Raedt and J. Ramon, Top-down induction of clustering trees, in J. W.
23 Shavlik (eds.), *Proceedings of the 15th International Conference on Machine Learning*
(Morgan Kaufmann, 1998), pp. 55–63.
- 24 [4] S. Bloehdorn and Y. Sure, Kernel methods for mining instance data in ontologies, in
25 K. Aberer *et al.*, (eds.), *Proceedings of the 6th International Semantic Web Conference*,
26 LNCS Vol. 4825, 2007, pp. 58–71.
- 27 [5] C. d’Amato, S. Staab and N. Fanizzi, On the influence of description logics ontologies
28 on conceptual similarity, in A. Gangemi and J. Euzenat (eds.), *Proceedings of 16th*
29 *International Conference on Knowledge Engineering: Practice and Patterns*,
EKA W2008, LNCS Vol. 5268, 2008, pp. 48–63.
- 30 [6] N. Fanizzi, Concept induction in description logics using information-theoretic heur-
31 istics, *Int. J. Semantic Web Inf. Syst.* **7**(2) (2011) 23–44.
- 32 [7] N. Fanizzi, C. d’Amato and F. Esposito, Induction of concepts in web ontologies through
33 terminological decision trees, in José L. Balcázar *et al.* (eds.), *Proceedings of European*
34 *Conference on Machine Learning and Knowledge Discovery in Databases, ECML/*
PKDD2010, Part I, LNAI Vol. 6321, 2010, pp. 442–457.
- 35 [8] N. Fanizzi, C. d’Amato and F. Esposito, Learning to rank individuals in description
36 logics using kernel perceptrons, in P. Hitzler and T. Lukasiewicz (eds.), *Proceedings of*
37 *4th International Conference on Web Reasoning and Rule Systems*, LNCS Vol. 6333
(Springer 2010), pp. 173–181.
- 38 [9] Ramanathan V. Guha and D. Brickley, RDF vocabulary description language 1.0: RDF
39 schema, W3C recommendation, W3C, February 2004, [http://www.w3.org/TR/2004/](http://www.w3.org/TR/2004/REC-rdf-schema-20040210/)
40 [REC-rdf-schema-20040210/](http://www.w3.org/TR/2004/REC-rdf-schema-20040210/).
- 41 [10] J. Han, *Data Mining: Concepts and Techniques* (Morgan Kaufmann, 2005).
- 42 [11] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning – Data*
43 *Mining, Inference, and Prediction* (Springer, 2001).

18 *N. Fanizzi et al.*

- 1 [12] S. Hellmann, J. Lehmann and S. Auer, Learning of OWL class descriptions on very large
2 knowledge bases, *Int. J. Semantic Web Inf. Syst.* **5**(2) (2009) 25–48.
- 3 [13] L. Iannone, I. Palmisano and N. Fanizzi, An algorithm based on counterfactuals for
4 concept learning in the semantic web, *Applied Intelligence* **26**(2) (2007) 139–159.
- 5 [14] A. Karalic and I. Bratko, First order regression, *Machine Learning* **26** (1997) 147–176.
- 6 [15] S. Kramer, Structural regression trees, in *Proceedings of 13th National Conference on*
7 *Artificial Intelligence* (AAAI Press/MIT Press) 1996, pp. 812–819.
- 8 [16] S. Kramer and G. Widmer, Inducing classification and regression trees in first order logic, in
9 S. Džeroski and N. Lavrač (eds.), *Relational Data Mining* (Springer, 2001), pp. 140–156.
- 10 [17] J. Lehmann and P. Hitzler, Concept learning in description logics using refinement
11 operators, *Machine Learning* **78**(1–2) (2010) 203–250.
- 12 [18] L. Torgo, Functional models for regression tree leaves, in D. H. Fisher (eds.), *Proceedings*
13 *of 14th International Conference on Machine Learning* (Morgan Kaufmann, 1997),
14 pp. 385–393.
- 15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43