

Prediction of Class and Property Assertions on OWL Ontologies through Evidence Combination

Giuseppe Rizzo
Computer Science Dept.
University of Bari, Italy
rizzo.giuseppe@unibari.net

Claudia d'Amato
Computer Science Dept.
University of Bari, Italy
claudia.damato@di.uniba.it

Nicola Fanizzi
Computer Science Dept.
University of Bari, Italy
fanizzi@di.uniba.it

Floriana Esposito
Computer Science Dept.
University of Bari, Italy
esposito@di.uniba.it

ABSTRACT

In the line of our investigation of inductive methods for Semantic Web reasoning, we propose an alternative way for approximate ABox reasoning based on the evidence and the analogical principle of the nearest-neighbors. Once neighbors of a test individual are selected through some distance measures, a combination rule descending from the Dempster-Shafer theory can join together the evidence provided by the various neighbor individuals in order to predict unknown values in a learning problem. We show how to exploit the procedure in the problems of determining unknown class- and role-memberships or fillers for datatype properties which may be the basis for many further ABox inductive reasoning algorithms.

Keywords

Similarity, Distances, Evidence Theory, Belief, Plausibility, Combination Rule, Description Logics

1. INTRODUCTION

In the context of reasoning in the Semantic Web, a growing interest is being shown to alternative procedures extending the standard methods so that they can deal with the various facets of uncertainty related with Web reasoning [19].

Extensions of the classic probability measures [18] offer alternative ways to deal with inherent uncertainty of the knowledge bases in the Semantic Web. Particularly, belief and plausibility measures adopted in the *Dempster-Shafer Theory of Evidence* [18] have been exploited as means for dealing with incompleteness [9] and also inconsistency [20], which may arise from the aggregation of data and metadata on a large and distributed scale. The Dempster-Shafer theory is a generalization of the Bayesian theory of subjective probability in which functions base degrees of belief on the

probability of a proposition. Belief functions represent the probability that a given proposition is provable from a set of other propositions, to which probabilities are assigned [21].

In this work we undertake again the inductive point of view. Indeed, in many Semantic Web domains a very large number of assertions can potentially be true but often only a small number of them is known to be true or can be inferred to be true. So far the application of combination rules related to the Dempster-Shafer theory in this field has regarded the induction of metrics which are essential for all similarity-based reasoning methods [9]. One of the applications of such measures was related to the prediction of assertions through nearest neighbor procedures. A general-purpose evidential nearest neighbor procedure based on the Dempster-Shafer combination rule (a generalization of the classic Bayes rule) has been proposed [7]. In this work this method is extended to the specific case of semantic knowledge bases through a more epistemically appropriate combination procedure [25], exploiting specific metrics to assess the similarity of the individuals involved.

In the perspective of inductive methods, the need for a definition of a semantic similarity measure for *individuals* arises, that is a problem that so far received less attention in the literature compared to the measures for concepts. Recently proposed dissimilarity measures for individuals in specific languages founded in *Description Logics* [1] turned out to be practically effective for the targeted inductive tasks [4], however they are still based on structural criteria so that they can hardly scale to more complex languages. We devised families of dissimilarity measures for semantically annotated resources, which can overcome the aforementioned limitations [5, 12]. Our measures are mainly based on the Minkowski's norms for Euclidean spaces induced by means of a method developed in the context of relational *machine learning* [15]. Namely, the measures are based on the degree of discernibility of the input individuals with respect to a given context [13] (or committee of features), which are represented by concept descriptions expressed in the language of choice.

The main contributions of this are regarding the extension of a framework for the classification of individuals through a prediction procedure based on evidence theory and similarity [11]. In particular we propose using Yager's rule of combination and exploiting the mentioned families of metrics defined for individuals in ontologies. This allows for mea-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WIMS'11 May 25-27, 2011 Sogndal, Norway.

Copyright 2011 ACM 978-1-4503-0148-0/11/05 ...\$10.00.

asuring the confirmation of the truth of candidate assertions. The prediction of the values (related to class-membership or datatype and object properties) may have plenty of applications in uncertainty reasoning with ontologies.

The remainder of the paper is organized as follows. In the next section (§2), we recall the definition of distance measures that shall be utilized for selecting neighbor individuals. Then (§3), the basics of the Dempster-Shafer theory and a nearest-neighbor procedure based on an alternative rule of combination are recalled. Hence (§4) we present the applications of the method to the problems of determining the class- or role-membership of individuals w.r.t. given query concepts / roles as well as the prediction of fillers for datatype properties. An experimental evaluation of the method in three prediction tasks is showed in (§5). Relevant related work are discussed in (§6) and we conclude (§7) by proposing extensions and applications of these methods in further works.

2. DISSIMILARITY MEASURES BETWEEN INDIVIDUALS

Since the reasoning method to be presented in the following is intended to be general purpose, no specific language, will be assumed in the following for resources, concepts (classes) and their properties. It suffices to consider a generic representation that can be mapped to some Description Logic language with the standard model-theoretic semantics (see [1] for a thorough reference).

A *knowledge base* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ comprises a *TBox* $\mathcal{T}$ and an *ABox* $\mathcal{A}$. $\mathcal{T}$ is a set of axioms concerning the (partial) definition of concepts (and roles) through class (role) expressions. $\mathcal{A}$ contains assertions (ground facts) concerning the world state. The set of the individuals occurring in $\mathcal{A}$ will be denoted with $\text{Ind}(\mathcal{A})$. Each individual can be assumed to be identified by its own URI (it is useful in this context to make the *unique names assumption*).

Similarity-based tasks, such as individual classification, retrieval, and clustering require language-independent measures for individuals whose definition can capture semantic aspects of their occurrence in the knowledge base [5, 12].

For our purposes, we need functions to assess the similarity of individuals. However individuals do not have an explicit syntactic (or algebraic) structure that can be compared (unless one resorts to language-specific notions [4], such as the *most specific concept* [1]). Hence it turns out to be difficult to adapt measures for concepts that have been recently proposed, e.g. see [3, 8, 17].

Focusing on the semantic level, the leading idea may be that, similar individuals should behave similarly with respect to the same concepts. A way for assessing the similarity of individuals in a knowledge base can be based on the comparison of their semantics along a number of dimensions represented by a set of concept descriptions (henceforth referred to as the *committee* or *context* [13]). Specifically, the measure may compare individuals on the grounds of their behavior w.r.t. a given context, say $\mathbf{C} = \{C_1, C_2, \dots, C_m\}$, which stands as a group of discriminating relevant concepts (*features*) expressed in the considered language.

We begin with defining the behavior of an individual w.r.t. a certain concept in terms of projecting it in this dimension: Given a concept $C_i \in \mathbf{C}$, the related *projection function* π_i :

$\text{Ind}(\mathcal{A}) \mapsto [0, 1]$ is defined:

$$\forall a \in \text{Ind}(\mathcal{A}) \quad \pi_i(a) = \begin{cases} 1 & \mathcal{K} \models C_i(a) \\ 0 & \mathcal{K} \models \neg C_i(a) \\ \pi_i & \text{otherwise} \end{cases}$$

The intermediate value π_i corresponds to the case when a reasoner cannot assign the truth value for a certain membership query. This is due to the *Open World Assumption* normally made in Semantic Web reasoning. Hence, as in the classic probabilistic models, the prior membership probability π_i w.r.t. F_i is considered, if known. Otherwise a uniform distribution is assumed ($\pi_i = \frac{1}{2}$). Further ways to approximate these values in case of uncertainty are investigated in [9]. For example the ratio of the concept extension w.r.t. the domain of objects has been approximated by the ratio between the cardinality of the retrieval of the concept and the number of individuals occurring in the knowledge base $|\text{Ind}(\mathcal{A})|$.

The discernibility functions compare two input individuals w.r.t. the concepts in the context through their projections: Given a feature concept $C_i \in \mathbf{C}$, the related *discernibility function* $\delta_i : \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \mapsto [0, 1]$ is defined as follows:

$$\forall (a, b) \in \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \quad \delta_i(a, b) = |\pi_i(a) - \pi_i(b)|$$

Finally, a family of dissimilarity measures for individuals inspired to the Minkowski's metrics can be defined as follows [5, 12]: Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a knowledge base. Given a context $\mathbf{C}$ and a related vector of weights $\vec{w}$, a family of dissimilarity measures $\{d_p^{\mathbf{C}}\}_{p \in \mathbf{N}}$, is made up of the functions $d_p^{\mathbf{C}} : \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \mapsto [0, 1]$ defined as follows:

$$\forall (a, b) \in \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \quad d_p^{\mathbf{C}}(a, b) = \left[\sum_{C_i \in \mathbf{C}} w_i \delta_i(a, b)^p \right]^{\frac{1}{p}}$$

The effect of the weights is to normalize w.r.t. the other features involved. Obviously these measures are not absolute, then they should be also be considered w.r.t. the context of choice, hence comparisons across different contexts may not be meaningful. Larger contexts are likely to decrease the measures because of the normalizing factor yet these values is affected also by the degree of redundancy of the features employed. In other works the choice of the weights is done according to variance or entropy associated to the various concepts in the context [5, 12].

Compared to other proposed measures [3, 4, 6], the presented functions do not depend on the constructors of a specific language, rather they require only (retrieval or) instance-checking inferences for computing the projections through class-membership queries to the knowledge base.

The complexity of measuring the dissimilarity of two individuals depends on the complexity of such inferences (see [1], Ch. 3). Note also that the projections that determine the measure can be computed (or derived from statistics maintained on the knowledge base) before the actual distance application, thus determining a speed-up in the computation of the measure. This is very important for algorithms that massively use this distance, such as all instance-based methods.

One should assume that $\mathbf{C}$ represents a set of (possibly redundant) features that are able to discriminate individuals that are actually different. The choice of the concepts to be included (a *feature selection* problem [15]) may be

crucial. Therefore, specific optimization algorithms founded in *randomized search* have been devised which are able to find optimal choices of discriminating contexts [5, 12]. However, the results obtained so far with knowledge bases drawn from ontology libraries showed that (a selection) of the primitive and defined concepts are often sufficient to induce sufficiently discriminating measures.

3. EVIDENCE-THEORETIC PREDICTION

In this section the basics of the theory of evidence and combination rules [18] are recalled. Then a nearest neighbor classification procedure based on the rule of combination [7] is extended in order to perform prediction of unobserved values (related to datatype properties or also class-membership).

3.1 Basics of the Evidence Theory

In the Dempster-Shafer theory, a *frame of discernment* Ω is defined as the set of all hypotheses in a certain domain. Particularly, in a classification problem it is the set of all possible classes. A *basic belief assignment* (BBA) is a function m that defines a mapping $m : 2^\Omega \mapsto [0, 1]$ verifying: $\sum_{A \in \Omega} m(A) = 1$. Given a certain piece of evidence, the value of the BBA for a given set A expresses a measure of belief that is committed exactly to A . The quantity $m(A)$ pertains only to A and does not imply any additional claims about any of its subsets. If $m(A) > 0$, then A is called a *focal element* for m .

The BBA m cannot be considered a proper probability measure: it is defined over 2^Ω instead of Ω and it does not require the properties of monotone measures [18]. The BBA m and its associated focal elements define a *body of evidence*, from which a *belief function* Bel and a *plausibility function* Pl can be derived as mappings from 2^Ω to $[0, 1]$. For a given $A \subseteq \Omega$, the *belief* in A , denoted $Bel(A)$, represents a measure of the total belief committed to A given the available evidence. Bel is defined as follows:

$$\forall A \in 2^\Omega \quad Bel(A) = \sum_{\emptyset \neq B \subseteq A} m(B) \in [0, 1] \quad (1)$$

Analogously, the *plausibility* of A , denoted $Pl(A)$, represents the amount of belief that could be placed in A , if further information became available. Pl is defined as follows:

$$\forall A \in 2^\Omega \quad Pl(A) = \sum_{B \cap A \neq \emptyset} m(B) \in [0, 1] \quad (2)$$

It is easy to see that: $Pl(A) = Bel(\Omega) - Bel(\bar{A})$. Moreover $m(\emptyset) = 1 - Bel(\Omega)$ and for each $A \neq \emptyset$: $m(A) = \sum_{B \subseteq A} (-1)^{|A \setminus B|} Bel(B)$. Using these equations, knowing just one function among m , Bel , and Pl allows to derive the others.

The Dempster-Shafer rule of combination [18] is an operation for pooling evidence from a variety of sources. This rule aggregates independent bodies of evidence defined within the same frame of discernment into one body of evidence. Let m_1 and m_2 be two BBAs. The new BBA obtained by combining m_1 and m_2 using the rule of combination, m_{12} is the orthogonal sum of m_1 and m_2 . Generally, the normalized version of the rule is used:

$$\forall A \in 2^\Omega \setminus \{\emptyset\}$$

$$m_{12}(A) = (m_1 \oplus m_2)(A) = \frac{\sum_{B \cap C = A} m_1(B) m_2(C)}{1 - \sum_{B \cap C = \emptyset} m_1(B) m_2(C)}$$

(and $m_{12}(\emptyset) = 0$) where the numerator $(1 - c)$ normalizes the values of the combined BBA w.r.t. the amount of conflict c between m_1 and m_2 .

Different evidence fusion rules have been proposed [18]. A more epistemologically sound combination rule [25] for our purposes places the probability mass related to the conflict between the BBAs to the case of maximal ignorance:

$$\forall A \in 2^\Omega$$

$$m_{12}(A) = \begin{cases} \frac{\sum_{B \cap C = A} m_1(B) m_2(C)}{m_1(\Omega) m_2(\Omega) + c} & A \neq \Omega \wedge A \neq \emptyset \\ 0 & A = \Omega \\ 0 & A = \emptyset \end{cases}$$

This means that the conflict between the two sources of evidence is not hidden, but it is explicitly recognized as a contributor to ignorance.

Due to the associativity and commutativity of the operations involved, it is easy to prove that the resulting combination operator is associative and commutative, and admits the vacuous BBA (Ω unique focal set) as neutral element.

3.2 The Evidential Nearest-Neighbors Procedure

Let us consider the finite set of instances X and a finite set of integers $V \subseteq \mathbf{Z}$ to be used as labels (which may correspond to disjoint classes or distinct attribute values). The available information is assumed to consist in a training set $\text{TrSet} = \{(x_1, v_1), \dots, (x_M, v_M)\} \subseteq \text{Ind} \times V$ of single-labeled instances (*examples*). In our case, $X = \text{Ind}(A)$, the set of individual names occurring in the ontology.

Let x_q be a new individual to be classified on the basis of its nearest neighbors in TrSet . Let $N_k(x_q) = \{(x_{o(j)}, v_{o(j)}) \mid j = 1, \dots, k\}$ be the set of the k nearest neighbors of x_q in TrSet sorted by a function $o(\cdot)$ depending on an appropriate metric d which can be applied to ontology individuals (e.g., one of the measures in the family defined in the previous section §2).

Each pair $(x_i, v_i) \in N_k(x_q)$ constitutes a distinct item of evidence regarding the value to be predicted for x_q . If x_q is close to x_i according to d , then one will be inclined to believe that both instances are associated to the same value, while when $d(x_q, x_i)$ increases, this belief decreases and that leads to a situation of almost complete ignorance concerning the value to be predicted for x_q .

Consequently, each $(x_i, v_i) \in N_k(x_q)$ may induce a BBA m_i over V which can be defined as follows [7]:

$$\forall A \in 2^V$$

$$m_i(A) = \begin{cases} \lambda \sigma(d(x_q, x_i)) & A = \{v_i\} \\ 1 - \lambda \sigma(d(x_q, x_i)) & A = V \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where $\lambda \in]0, 1[$ is a parameter and $\sigma(\cdot)$ is a decreasing function such that $\sigma(0) = 1$ and $\lim_{d \rightarrow \infty} \sigma(d) = 0$ (e.g., $\sigma(d) = \exp(-\gamma d^n)$ with $\gamma > 0$ and $n \in \mathbf{N}$). The values of the parameters can be determined heuristically.

Considering each training individual in $N_k(x_q)$ as a separate source of evidence, k BBAs m_j are obtained. These can be pooled by means of the rule of combination leading to the aggregated BBA $\bar{m}$ that synthesizes the final belief:

$$\bar{m} = \bigoplus_{j=1}^k m_j = m_1 \oplus \dots \oplus m_k \quad (4)$$

In order to predict a value, functions $\bar{Bel}$ and $\bar{Pl}$ can be de-

rived from $\bar{m}$ using the equations seen above, and the query individual x_q is assigned the value in V that maximizes the belief or plausibility:

$$v_q = \operatorname{argmax}_{(x_i, v_i) \in N_k(x_q)} \overline{Bel}(\{v_i\})$$

or

$$v_q = \operatorname{argmax}_{(x_i, v_i) \in N_k(x_q)} \overline{Pl}(\{v_i\})$$

The former choice (select the hypothesis with the greatest degree of belief i.e. the most credible) corresponds to a *skeptical* viewpoint while the latter (select the hypothesis with the lowest degree of doubt i.e. the most plausible) is more *credulous*. The degree belief (or plausibility) of the predicted value provides also a way to compare the answers of an algorithm built on top of such analogical procedure. This is useful for tasks such as ranking, matchmaking, etc..

Finally, it is possible to combine the two measures Bel and Pl analogously to necessity (Nec) and possibility (Pos) in *Possibility Theory* (which can be considered a special case¹ of Dempster-Shafer theory). One can define a single measure of *confirmation* C , ranging in $[-1, +1]$, by means of a simple one-to-one transformation [18]:

$$\forall A \subseteq \Omega \quad C(A) = Bel(A) + Pl(A) - 1 \quad (5)$$

Hence, denoted with $\overline{C}$ the combination of $\overline{Bel}$ and $\overline{Pl}$, the resulting rule for predicting the uncertain value for the test individual can be written as follows:

$$v_q = \operatorname{argmax}_{(x_i, v_i) \in N_k(x_q)} \overline{C}(\{v_i\}) \quad (6)$$

Summing up, the procedure is as reported as Algorithm 1:

Algorithm 1 The evidence nearest neighbor procedure.

$ENN_k(x_q, \text{TrSet}, V)$

Input:

x_q : query individual, TrSet : training set, V value set

Output:

$v \in V$: predicted value

- 1: Compute the neighbor set $N_k(x_q) \subseteq \text{TrSet}$.
 - 2: **for all** $i \leftarrow 1$ **to** k **do**
 - 3: Compute m_i (Eq. 3)
 - 4: **end for**
 - 5: **for all** $v \in V$ **do**
 - 6: Compute $\bar{m}$ (Eq. 4) and derive $\overline{Bel}$ and $\overline{Pl}$ (Eqs. 1–2)
 - 7: Compute the confirmation $\overline{C}$ (Eq. 5) from $\overline{Bel}$ and $\overline{Pl}$
 - 8: **end for**
 - 9: Select $v \in V$ that maximizes $\overline{C}$ (Eq. 6).
-

It is worthwhile to note that the complexity of the method is polynomial in the number of instances in the TrSet . If this set is compact and contains very prototypical individuals with plenty of related assertions, then the resulting predictions are likely to be accurate. Another source of complexity in the computations may be the number of values in V which may yield a large number of subsets $2^{|V|}$ for which BBAs are to be computed. However this depends also on the kind of problem that is to be solved (e.g., in class membership detection $|V| = 2$). Moreover what really matters in the number

¹Precisely, the body of evidence must contain *consonant* focal sets, i.e. when the set of focal elements is a nested family [18].

of focal sets for each BBA which may be much less than $2^{|V|}$.

4. ASSERTION PREDICTION

The utility of the presented procedure when applied to ontology reasoning can be manifold. In the following we propose its employment in the inductive prediction of unknown values related to class-membership and datatype / object property fillers. This feature may be easily embedded in an ontology management system in order to help the knowledge engineers elicit assertions which may be not be derived from the knowledge base yet they can be rather made in analogy with the others [4].

In the following, the symbol $\approx$ in expressions like $\mathcal{K} \approx \alpha$ will denote the derivation of the assertion α from the knowledge base $\mathcal{K}$ obtained through an alternative procedure (like the evidence nearest neighbor presented in the previous section).

4.1 Class-Membership

Let us suppose a (query) concept Q is given. In this case one may consider only examples made up of individuals with a definite class-membership leading to a binary problem with a set of values $V_Q = \{+1, -1\}$ denoting, resp., membership and non-membership w.r.t. the query concept. Alternatively, one may admit ternary problems with some labels set to 0 to explicitly denote an indefinite (uncertain) class-membership [4, 5]. We shall also consider the related training set $\text{TrSet}_Q \subseteq \text{Ind}(\mathcal{A}) \times V_Q$. The values of the labels v_i for the training examples can be obtained through deductive reasoning (instance-checking) or specific facilities made available by the knowledge management systems [16].

Now in order to predict the class-membership value v_q for some individual x_q w.r.t. Q , it suffices to call the procedure $ENN_k(x_q, \text{TrSet}_Q, V_Q)$ and decide on the grounds of the returned value. Thus in a binary setting ($V_Q = \{+1, -1\}$), one will either conclude that $\mathcal{K} \approx Q(x_q)$ or $\mathcal{K} \approx \neg Q(x_q)$ depending on the value that maximizes $\overline{C}$ in Eq. 6 (resp., $v_q = +1$ or $v_q = -1$). Moreover the value of the confirmation function which determined the returned value v_q can be exploited for ranking the hits by comparing the strength of the inductive conclusions.

Adopting a ternary setting, it may turn out that the most likely value is $v_q = 0$ resulting in an uncertain case. One may force the choice among the values of $\overline{C}$ for $v_q = -1$ and $v_q = +1$ e.g., when the confirmation degree exceeds some given threshold.

The inductive procedure described above can be trivially exploited for performing the retrieval of a certain concept inductively. Given a certain concept Q , it would suffice to find all individuals $a \in \text{Ind}(\mathcal{A})$ that are such that $\mathcal{K} \approx Q(a)$. The hits could be returned ranked by the respective confirmation value $\overline{C}(+1)$.

4.2 Datatype Fillers

In this case, let us suppose a certain (functional) datatype property P is given and the problem is to predict its value for a certain test individual a (which has to be supposed to belong to the domain of the property). The set of values V_P may correspond to the (discrete and finite) range of the property or to its restriction to the observed values for the training instances: $V_P = \{v \in \text{range}(P) \mid \exists P(a, v) \in \mathcal{A}\}$. Different settings may be devised allowing for some special

value(s) denoting the case of a yet unobserved value(s) for that property.

The related training set will be some $\text{TrSet}_P \subseteq \text{domain}(P) \times V_P$, where $\text{domain}(P) \subseteq \text{Ind}(\mathcal{A})$ is the set of individual names that have a known P -value in the knowledge base. Differently from the previous problem, datatype properties generally do not have a specific intensional definition in the knowledge base (except for the specification of domain and range), hence a mere look-up in the ABox should suffice to determine the TrSet .

Now to predict the value in V_P of the datatype property P for some individual a , the method requires calling the procedure with $\text{ENN}_k(a, \text{TrSet}_P, V_P)$. Thus in this setting, if v_q is the value that maximizes Eq. 6 then we can write $\mathcal{K} \approx P(a, v_q)$. Also in this case the value of the confirmation function which determined choice of the value v_q can be exploited for comparing the strength of an inductive conclusion to others.

In case of special settings with dummy values indicating unobserved values, when these are found to be the most credible among the others, a knowledge engineer should be contacted for the necessary changes to the ontology.

The inductive procedure described above can be trivially exploited for performing alternate forms of retrieval e.g., finding all individuals with a certain value for the given property. Given a certain value v , it would suffice to find all individuals $a \in \text{Ind}(\mathcal{A})$ that are such that $\mathcal{K} \approx P(a, v)$. Again, the hits could be returned ranked according to the respective confirmation value $\bar{C}(+1)$.

The limitation of treating only functional datatype properties may be overcome by considering a different way to assign the probability mass to BBAs than Eq. 3, including sub-sets of all possible values. Examples are to be constructed accordingly (labels will be chosen in 2^{V_P}). Alternatively, more complex frames of discernment e.g., $\Omega' = 2^\Omega$, so consider sets of values as possible fillers of the property. In all such settings the computation of the BBAs and descending measures would become of course much more complex and expensive, yet clever solutions (or approximations) proposed in the literature [7] may contribute to mitigate this problem.

4.3 Relationships among Individuals

In principle, a very similar setting may be used in order to establish the possibility that a certain test individual is related through some object property with some other individual.

Since the set $\text{Ind}(\mathcal{A})$ is finite (the target is not discovering relations with unseen individuals), one may want to find all individuals that are related to a test one through some object property, say R . The problem can be decomposed into smaller ones aiming at verifying whether $\mathcal{K} \approx R(a, b)$ holds, as shown in Algorithm 2.

Note that, in the construction of the training sets, the inference $\mathcal{K} \models R(x, b)$ may turn out to be merely an ABox lookup operation for the given assertions (when roles are not intensionally defined in a proper RBox). Conversely, if an RBox is available (sometimes as a subset of the TBox) the values of the label for the training examples can be obtained through deductive reasoning (instance-checking) or the mentioned facilities made available by advanced reasoners or knowledge management systems [16].

This simple setting makes a sort of *closed-world assumption* in the decision of the induced assertions descending

Algorithm 2 Finding relationships among individuals.

REL_PREDICTION($\text{Ind}(\mathcal{A})$)

Input: $\text{Ind}(\mathcal{A})$ set of individuals

Output: $\approx$ for the individuals in $\text{Ind}(\mathcal{A})$

```

1: for all  $b \in \text{Ind}(\mathcal{A})$  do
2:   for all  $a \in \text{Ind}(\mathcal{A})$  do
3:      $\text{TrSet} \leftarrow \{(x, v) \mid x \in \text{Ind}(\mathcal{A}) \setminus \{a\},$ 
         $\text{if } \mathcal{K} \models R(x, b) \text{ then } v \leftarrow +1 \text{ else } v \leftarrow -1\}$ 
4:      $v_b^R \leftarrow \text{ENN}_k(a, \text{TrSet}, \{+1, -1\})$ 
5:     if  $v_b^R = +1$  then
6:       Predict  $\mathcal{K} \approx R(a, b)$ 
7:     else
8:       Predict  $\mathcal{K} \not\approx R(a, b)$ 
9:     end if
10:  end for
11: end for
```

from the adoption of the binary value set and the composition of the TrSet . A more cautious setting would involve a ternary value set $V_R = \{-1, 0, +1\}$ which allows for an explicit treatment of those individuals a for which $R(a, b)$ is not derivable (or just absent from the ABox). The final decision on the induced conclusion has to consider also this new possibility (e.g., using a threshold of confirmation for accepting likely assertions).

5. EVALUATION

5.1 Experiment Design

In order to test the algorithms on real ontologies, the resulting system prototype was applied to the three kinds of prediction problems involving the individuals therein.

To this purpose, a number of ontologies from different domains represented in OWL have been selected², namely: FINITESTATEMACHINES (FSM) concerning finite state machines, NEWTESTAMENTNAMES (NTN) accounting for characters and places mentioned in that book, the *Internet Movie DataBase* ontology (IMDB), the BioPax glycolysis ontology (BioPax) describing the glycolysis pathway from the EcoCyc database, translated into BioPax format, and two medical knowledge bases, the *Breast Clinical Ontology* (BCO) and *Human Diseases* HDIs. Tab. 1 summarizes important details concerning these ontologies, in terms of the numbers of concepts, object and datatype properties and individuals.

The experiments have been replicated adopting four different rules of evidence of combination, namely: Dempster-Shafer's, Dubois-Pradé's, Yager's and the Mixing rule [18]. A 10-fold cross validation design was adopted to determine the average performance indices (described below). Given the training set TS in each experiment run, the *size of the neighborhood* was set to $k = \log |TS|$. As regards the settings for the other parameters $\lambda = .95$, and the other are determined heuristically (see [7]): n is a small integer (typically $n = 2$) and has little impact on the function applied to the distance among examples, while we set $\gamma = 1/d^*$, where d^* is the average distance among training instances belonging to the target class. All atomic concepts defined in each

²The ontologies can be found in standard repositories: the Protégé library (<http://protege.stanford.edu/plugins/owl/owl-library>) and TONES (<http://owl.cs.manchester.ac.uk/repository>).

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

Ontology	DL language	#concepts	#object properties	#datatype properties	#individuals
FSM	$\mathcal{SOF}(\mathcal{D})$	20	10	7	37
BCO	$\mathcal{ALCROF}(\mathcal{D})$	196	22	3	112
IMDB	$\mathcal{ALIN}(\mathcal{D})$	7	5	13	302
BioPAX	$\mathcal{ALCIF}(\mathcal{D})$	74	70	40	323
HDIs	$\mathcal{ALCIF}(\mathcal{D})$	1498	10	15	639

ontology have been considered to be included in the set of features for the contexts C .

The experiments concerning the three specific tasks required suitable ontologies to create suitable training and test sets for the task. As in previous experiments on the same task (e.g. see [5, 10]) random classification problems have been generated by constructing concept descriptions using the concepts and relations offered by each ontology and the operators of the $\mathcal{ALC}$ languages. The role-filling tasks have been tested by randomly selecting 5 properties from the ontologies. Finally, only ontologies containing functional properties were usable for the the problem of predicting the values of datatype properties for given individuals. Besides only ranges with small cardinalities are currently tractable by the system prototype.

5.2 Outcomes

Due to the open-world semantics, it may happen that the membership of an individual w.r.t. a query cannot be determined by a reasoner, since it can build models for the membership w.r.t. both the concept $C(a)$ and its negation $\neg C(a)$. Then a three-way classification was adopted and evaluated using the following indices already adopted in previous works [5, 9, 10]. Essentially they measure the correspondence between the deductive and inductive classification for the instances w.r.t. the query concept provided, resp., by the reasoner and the inductive algorithm:

- *match rate* (M%), i.e. number of cases of individuals that got the same classification with both modes;
- *omission error rate* (O%), amount of individuals for which the membership w.r.t. the given query could not be determined using the inductive method, while they can be proven to belong to the query concept or to its complement;
- *commission error rate* (C%), amount of individuals found to belong to the query concept according to the inductive classification, while they can be proven to belong to its complement and vice-versa;
- *induction rate* (I%), amount of individuals found to belong to the query concept or its complement according to the inductive classifier, while either case is not logically derivable from the knowledge base.

For each index, average value and standard deviation over the various folds is reported in the following tables.

Table 2 presents the outcomes of the experiments where the method aimed at determining the class membership of the test individuals. The table shows a good performance of the inductive classification in terms of the average match rates. it also shows that omission errors occurred quite

rarely, while commission rates are slightly higher (especially with the smallest ontology). Finally the average induction rates are sensible (see especially the outcomes of the BCO case). This means that the system can actually help determining a class-membership by analogy when the reasoner cannot.

Comparing the outcomes in terms of the various evidence combination rules, no significant differences were observed not in terms of average values and of their variance (standard deviation). These outcomes are on average in line with those obtained through more sophisticated such as those methods exploiting kernel methods (e.g. SVMs in [10]).

Table 3 presents the outcomes of the experiments where the method aimed at predicting the individuals that are likely to fill a given property involving the test individuals. The table shows an optimal performance of the inductive method in terms of the average match rates. Consequently the other cases (omission and commission error, induction) were not so frequent to be statistically significant. Also in terms of the evidence combination rules, no significant difference was observed. The careful consideration of the ontologies shows that they are likely too easy for the task for the primitive roles in the ontology where used. Probably having a chance of combining different roles would provide harder learning problems.

Finally, Table 4 shows the results of the experiments where the method aimed at determining the value of a functional datatype property for the test individuals. Preliminarily, note that the induction rate is null also because the considered properties had a value for almost all of the considered individuals. Similarly, omission errors occurred quite rarely. The table shows a fair performance of the inductive classification in terms of the average match rates. However the commission rates are considerably high. Another noticeable feature is the higher variance w.r.t. the outcomes of the previous experiments. Comparing the outcomes in terms of the various evidence combination rules, while they coincide for the case of BCO, in the other cases sensible differences were observed in terms of average values and of their variance (standard deviation). Specifically, Yager’s rule seemed a little better, while with Dubois-Prade’s rule the performance decreased.

6. RELATED WORK

The proposed method is related to those approaches devised to offer alternative ways of reasoning with ABoxes for eliciting hidden knowledge (regularities) in order to complete and populate the ontology with likely assertions even in the occurrence of incorrect parts, supposing this kind of *noise* is not systematic.

The tasks of ontology completion and population have often been tackled through formal methods (such as *formal*

Table 2: Outcomes of the experiments of class-membership prediction with the four combination rules: average values of the indices and standard deviations.

Ontology		Dempster	Dubois-Prade	Mixing	Yager
FSM	M%	86.60 ± 04.42	84.75 ± 04.49	85.80 ± 03.90	89.00 ± 04.65
	C%	04.69 ± 03.05	06.65 ± 03.06	05.49 ± 02.33	02.29 ± 02.76
	O%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	I%	08.71 ± 00.29	08.71 ± 00.29	08.71 ± 00.29	08.71 ± 00.29
BioPAX	M%	94.93 ± 00.32	94.76 ± 00.32	94.93 ± 00.32	94.93 ± 00.32
	C%	00.15 ± 00.00	00.32 ± 00.00	00.15 ± 00.00	00.15 ± 00.00
	O%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	I%	04.91 ± 00.29	04.91 ± 00.29	04.91 ± 00.29	04.91 ± 00.29
BCO	M%	85.21 ± 04.04	84.54 ± 04.83	85.21 ± 04.04	85.45 ± 04.18
	C%	00.81 ± 00.56	01.47 ± 01.54	00.81 ± 00.56	00.57 ± 00.70
	O%	00.05 ± 00.14	00.14 ± 00.23	00.05 ± 00.14	00.05 ± 00.14
	I%	13.93 ± 03.72	13.95 ± 03.64	13.93 ± 03.72	13.93 ± 03.72

Table 3: Outcomes of the experiments of object property filler prediction with the four combination rules: average values of the indices and standard deviations.

Ontology		Dempster	Dubois-Prade	Mixing	Yager
FSM	M%	99.64 ± 00.33	99.64 ± 00.33	99.98 ± 00.07	99.64 ± 00.33
	C%	00.02 ± 00.07	00.36 ± 00.33	00.02 ± 00.07	00.36 ± 00.33
	O%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
BioPAX	M%	100.00 ± 00.00	100.00 ± 00.00	100.00 ± 00.00	100.00 ± 00.00
	C%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	O%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
BCO	M%	100.00 ± 00.00	100.00 ± 00.00	100.00 ± 00.00	100.00 ± 00.00
	C%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	O%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00

concept analysis [2]). Discovering new assertions (and related probabilities in a classical setting) is another related task for eliciting hidden knowledge in the ontologies. In [24] a machine learning method is proposed to estimate the truth of statements by exploiting regularities in the data. In [23] another statistical learning method for OWL-DL ontologies is proposed, combining a latent relational graphical model with Description Logic inference in a modular fashion. The probability of unknown role-assertions can be inductively inferred and known concept-assertions can be analyzed by clustering individuals.

Similarity-based reasoning with ontologies is the primary aim of this work which follows a number of related methods founded on dissimilarity measures for individuals in knowledge bases expressed in Description Logics [4, 5]. Mostly, they adopt some alternate form of the classic Nearest-Neighbor learning scheme [15] in order to draw inductive conclusions that often cannot be deductively entailed by the knowledge bases.

Similar approaches based on lazy learning have been proposed that adopt generalized probability theories such as the Dempster-Shafer. In [26], which was a source of inspiration for this paper, the standard rule of combination is exploited in an evidence-theoretic classification procedure where labels were not assumed to be mutually exclusive. Rules of combination had been used in [9] in order to learn precise metrics to be exploited in a lazy learning setting like those mentioned above.

One of the most appreciated advantages of performing inductive ABox reasoning through these methods is that they can naturally handle inconsistent (and inherently incomplete) knowledge bases, especially when inconsistency is not systematic. In [20] a method for dealing with inconsistent ABoxes populated through information extraction is proposed: it constructs ad hoc belief networks for the conflicting parts in an ontology and adopts the Dempster-Shafer theory for assessing the confidence of the resulting assertions.

7. CONCLUDING REMARKS AND OUTLOOK

In the line of our investigation of inductive methods for Semantic Web reasoning, we have proposed an alternative way for approximate ABox reasoning based on the nearest-neighbors analogical principle. Once neighbors of a test individual are selected through some distance measures, a combination rule descending from the Dempster-Shafer theory can fuse the evidence provided by the various neighbor individuals. We have shown how to exploit the procedure for assertion prediction problems such as determining unknown class- or role-memberships as well as attribute-values which may be the basis for many ABox inductive reasoning algorithms. The method has been implemented so to allow an empirical evaluation on real ontologies.

Special settings to accommodate cases of uncertain or un-

Table 4: Outcomes of the experiments of datatype property filler prediction with the four combination rules: average values of the indices and standard deviations.

Ontology		Dempster	Dubois-Prade	Mixing	Yager
BCO	M%	64.15 ± 13.53	33.79 ± 11.64	63.52 ± 15.08	71.14 ± 10.00
	C%	35.85 ± 13.53	13.61 ± 10.52	36.48 ± 15.08	28.86 ± 10.00
	O%	00.00 ± 00.00	52.60 ± 15.95	00.00 ± 00.00	00.00 ± 00.00
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
IMDB	M%	65.60 ± 06.38	39.73 ± 14.19	66.25 ± 05.94	61.34 ± 08.28
	C%	30.74 ± 06.57	13.62 ± 10.52	30.09 ± 06.13	35.00 ± 09.78
	O%	03.66 ± 03.74	43.01 ± 19.99	03.66 ± 03.74	03.66 ± 03.74
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00
HDIS	M%	61.00 ± 19.15	61.00 ± 19.15	61.00 ± 19.15	61.00 ± 19.15
	C%	35.62 ± 17.32	35.62 ± 17.32	35.62 ± 17.32	35.62 ± 17.32
	O%	03.38 ± 04.94	03.38 ± 04.94	03.38 ± 04.94	03.38 ± 04.94
	I%	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00	00.00 ± 00.00

observed values are to be investigated. One promising extension of the method concerns the possibility of considering infinite sets of values V following the studies [14, 18]. This would allow dealing with domains where the total amount of values is unknown (also due to the inherent nature of the Semantic Web). Moreover the predicted values often need not to be exclusive. Hence the prediction procedure would require an extension towards the consideration of sets of values instead of singletons.

As necessity and possibility measures are related to the belief measures (see note 1 at page) a natural extension may be towards the possibilistic theory and its calculus which is, in general, different from the Dempster-Shafer theory and calculus. Further possible extensions concern all other monotone measures such as the Sugeno λ -measures [18]. The extension towards the Possibility Theory is interesting also because of its parallelism with modal logics [14] and possibilistic extensions of Description Logics [22].

8. REFERENCES

- [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [2] F. Baader, B. Ganter, B. Sertkaya, and U. Sattler. Completing description logic knowledge bases using formal concept analysis. In M. Veloso, editor, *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 230–235, Hyderabad, India, 2007.
- [3] A. Borgida, T. Walsh, and H. Hirsh. Towards measuring similarity in description logics. In I. Horrocks, U. Sattler, and F. Wolter, editors, *Working Notes of the International Description Logics Workshop*, volume 147 of *CEUR Workshop Proceedings*, Edinburgh, UK, 2005.
- [4] C. d’Amato, N. Fanizzi, and F. Esposito. Analogical reasoning in description logics. In P. da Costa et al., editors, *Uncertainty Reasoning for the Semantic Web I*, volume 5327 of *LNAI*, pages 336–354. Springer, 2008.
- [5] C. d’Amato, N. Fanizzi, and F. Esposito. Query answering and ontology population: An inductive approach. In S. Bechhofer et al., editors, *Proceedings of the 5th European Semantic Web Conference*, *ESWC2008*, volume 5021 of *LNCS*, pages 288–302. Springer, 2008.
- [6] C. d’Amato, S. Staab, and N. Fanizzi. On the influence of description logics ontologies on conceptual similarity. In A. Gangemi and J. Euzenat, editors, *Proceedings of the 16th Knowledge Engineering Conference, EKAW2008*, volume 5268 of *LNAI*, pages 48–63. Springer, 2008.
- [7] T. Denoeux. A k-nearest neighbor classification rule based on Dempster-Shafer theory. *IEEE Transactions on Systems, Man and Cybernetics*, 25(5):804–813, 1995.
- [8] J. Euzenat and P. Shvaiko. *Ontology Matching*. Springer, 2007.
- [9] N. Fanizzi, C. d’Amato, and F. Esposito. Approximate measures of semantic dissimilarity under uncertainty. In P. da Costa et al., editors, *Uncertainty Reasoning for the Semantic Web I*, volume 5327 of *LNAI*, pages 355–372. Springer, 2008.
- [10] N. Fanizzi, C. d’Amato, and F. Esposito. Statistical learning for inductive query answering on OWL ontologies. In A. Sheth et al., editors, *Proceedings of the 7th International Semantic Web Conference, ISWC2008*, volume 5318 of *LNCS*, pages 195–212. Springer, 2008.
- [11] N. Fanizzi, C. d’Amato, and F. Esposito. Evidential nearest-neighbors classification for inductive abox reasoning. In F. Bobillo et al., editors, *Proceedings of the 5th International Workshop on Uncertainty Reasoning for the Semantic Web, URSW 2009*, volume 527 of *CEUR Workshop Proceedings*, pages 27–38. CEUR-WS.org, 2009.
- [12] N. Fanizzi, C. d’Amato, and F. Esposito. Metric-based stochastic conceptual clustering for ontologies. *Information Systems*, 34(8):725–739, 2009.
- [13] R. Goldstone, D. Medin, and J. Halberstadt. Similarity in context. *Memory and Cognition*, 25(2):237–255, 1997.
- [14] D. Harmanec, G. Klir, and Z. Wang. Modal logic interpretation of Dempster-Shafer theory: An infinite case. *International Journal of Approximate Reasoning*, 14(2-3):81–93, 1996.
- [15] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning – Data Mining*, volume 5021 of *LNCS*, pages 288–302. Springer, 2008.

- Inference, and Prediction*. Springer, 2001.
- [16] I. Horrocks, L. Li, D. Turi, and S. Bechhofer. The Instance Store: DL reasoning with large numbers of individuals. In V. Haarslev and R. Möller, editors, *Proceedings of the 2004 Description Logic Workshop, DL 2004*, volume 104 of *CEUR Workshop Proceedings*, pages 31–40. CEUR, 2004.
 - [17] K. Janowicz and M. Wilkes. Sim-dla: A novel semantic similarity measure for description logics reducing inter-concept to inter-instance similarity. In L. Aroyo et al., editors, *Proceedings of the 6th European Semantic Web Conference, ESWC2009*, volume 5554 of *LNCIS*, pages 353–367. Springer, 2009.
 - [18] G. Klir. *Uncertainty and Information*. Wiley, 2006.
 - [19] K. Laskey, K. Laskey, P. Costa, M. Kokar, T. Martin, and T. Lukasiewicz. *Uncertainty Reasoning for the World Wide Web*. W3C Incubator Group, 2008. <http://www.w3.org/2005/Incubator/urw3/XGR-urw3-20080331/>.
 - [20] A. Nikolov, V. Uren, E. Motta, and A. de Roeck. Using the Dempster-Shafer theory of evidence to resolve ABox inconsistencies. In P. da Costa et al., editors, *Uncertainty Reasoning for the Semantic Web I*, volume 5327 of *LNAI*, pages 143–160. Springer, 2008.
 - [21] J. Pearl. On probability intervals. *International Journal of Approximate Reasoning*, 2(3):211–216, 1988.
 - [22] G. Qi, J. Pan, and Q. Ji. A possibilistic extension of description logics. In D. Calvanese et al., editors, *Working Notes of the 20th International Description Logics Workshop, DL2007*, volume 250 of *CEUR Workshop Proceedings*, pages 435–442, Bressanone, Italy, 2007.
 - [23] A. Rettinger and M. Nickles. Infinite hidden semantic models for learning with OWL DL. In C. d’Amato et al., editors, *Proceedings of 1st ESWC Workshop on Inductive Reasoning and Machine Learning for the Semantic Web, IRMLoS09*, volume 474 of *CEUR Workshop Proceedings*, Heraklion, Greece, 2009.
 - [24] V. Tresp, Y. Huang, M. Bundschuh, and A. Rettinger. Materializing and querying learned knowledge. In C. d’Amato et al., editors, *Proceedings of 1st ESWC Workshop on Inductive Reasoning and Machine Learning for the Semantic Web, IRMLoS09*, volume 474 of *CEUR Workshop Proceedings*, Heraklion, Greece, 2009.
 - [25] R. Yager. On the Dempster-Shafer framework and new combination rules. *Information Sciences*, 41:93–137, 1987.
 - [26] Z. Younes, F. Abdallah, and T. Denoeux. An evidence-theoretic k-nearest neighbor rule for multi-label classification. In L. Godo and A. Pugliese, editors, *Proceedings of the 3rd International Conference on Scalable Uncertainty Management, SUM2009*, volume 5785 of *LNAI*, pages 297–308. Springer, 2009.