

Converting Classifications into OWL Ontologies

Fausto Giunchiglia¹, Ilya Zaihrayeu¹, and Feroz Farazi^{1,2}

¹Department of Information Engineering and Computer Science
University of Trento, Italy

²Department of Computer Science and Engineering
University of Chittagong, Bangladesh
{fausto, ilya, farazi}@disi.unitn.it

Abstract. Classification schemes, such as the DMoZ web directory, provide a convenient and intuitive way for humans to access classified contents. While being easy to be dealt with for humans, classification schemes remain hard to be reasoned about by automated software agents. Among other things, this hardness is conditioned by the ambiguous nature of the natural language used to describe classification categories. In this paper we describe how classification schemes can be converted into OWL ontologies, thus enabling reasoning on them by Semantic Web applications. The proposed solution is based on a two phase approach in which category names are first encoded in a concept language and then, together with the structure of the classification scheme, are converted into an OWL ontology. We demonstrate the practical applicability of our approach by showing how the results of reasoning on these OWL ontologies can help improve the organization and use of web directories.

1 Introduction

A *classification scheme*, or a *classification* for short, is a rooted tree whose nodes are assigned natural language labels and are populated with a (possibly empty) set of documents. Since the invention of classification by Aristotle in the 4th century BC, classifications have been used (and are still used) pervasively to represent various kinds of human knowledge. For example, classifications have been used in libraries (DDC¹, LCC² and Colon classification³); in Personal Knowledge Management (favorites, personal e-mails and folder hierarchies); and, lately, on the Web (Amazon⁴, Google⁵, Yahoo⁶).

While classifications are heavily used to categorize web contents, the evolution of the web foresees a more formal structure which can serve this purpose

¹ See <http://www.tnrdrlib.bc.ca/dewey.html>.

² See <http://www.loc.gov/catdir/cpsolcc.html>.

³ See <http://www.iskoi.org/doc/colon.htm>.

⁴ See <http://www.amazon.com>.

⁵ See <http://www.google.com>.

⁶ See <http://www.yahoo.com>.

– *ontology*, defined in Computer Science as *a specification of a conceptualization* [10]. Ontologies are core artifacts of Semantic Web, an extension of the current Web, in which information is given formal semantics such that computers can use inference rules to conduct automated reasoning on pieces of this information [1]. The key factor which makes this possible is the fact that ontologies are expressed in a formal language, suitable for automated reasoning.

In this paper we bridge the gap between informal classifications and formal ontologies by describing an approach to encoding classification labels in a formal language such that, together with the structure of the classification scheme, they can be then converted into OWL [2] ontologies (more precisely, into lightweight ontologies, as described in [9]). In principle, the proposed approach allows for automated reasoning on classifications through reasoning on corresponding OWL ontologies. Moreover, the conversion is fully automated. Web directories can be encoded into OWL ontologies without user intervention. We demonstrate the practical applicability of our approach by showing how the results of reasoning on these OWL ontologies can help improve the organization and use of classification schemes. While encoding classifications into a formal language is not new, the main novelty of this paper consists of converting classifications into OWL ontologies, which demonstrates a proof of concept that classifications can be seamlessly integrated in the Semantic Web infrastructure. The fully automated algorithm described in this paper is also novel, as well the characterization of the expressivity of the formal language (i.e. OWL Lite, OWL DL, OWL Full) needed to encode classifications.

The rest of the paper is structured as follows. In Section 2 we describe a comparison between classification schemes and ontologies. In Section 3 we describe how to convert classification schemes into OWL ontologies and how the generated OWL ontologies can be enriched with additional axioms. In Section 4 we report the experimental results. Section 5 presents how this work helps in optimizing classifications. In Section 6 we discuss the related work and we conclude the paper in Section 7.

2 Classification Schemes vs Ontologies

In this section we discuss commonalities and differences between classifications and ontologies. In order to ground our discussion on well defined terms, below we give the definitions of these two kinds of artifacts.

A *classification* is a 5-tuple $C = \langle N, E, L, D, cl \rangle$ where N is a finite set of nodes, E is a set of edges on N , such that $\langle N, E \rangle$ is a rooted tree; L is a finite set of labels expressed in natural language, such that for any node $n_i \in N$, there is one and only one label $l_i \in L$; D is a set of documents and cl is a function which maps every $d_i \in D$ to a non-empty set of nodes $\{n_i\} \subseteq N$. In Figure 1 we show an example of a classification. Although classifications have no explicit formal semantics for edges, in this example we labeled each edge with the name of a hypothetical relation that may hold between the linked nodes.




Fig. 1. An example of a classification with link semantics made explicit.

An *ontology* is an *explicit specification of a conceptualization* [10]. They are often thought of as directed graphs whose nodes represent *concepts* and whose edges represent formal *relations* between concepts. The backbone structure of the ontology graph is a taxonomy in which all the relations are **sub-class-of**, whereas the remaining structure of the graph supplies auxiliary information about the modeled domain and may include relations like **part-of**, **located-in**, **is-parent-of**, and others [11]. Classes can be associated with instances through the **instance-of** relation. In Figure 2 we show an example of a small ontology.




Fig. 2. An example of an OWL ontology.

Even if both ontologies and classifications can often be represented in the form of a graph, ontologies and classifications are quite different in their uses, purpose, language, applications, and in other aspects which we summarize as follows:

- **Users:** a typical user of classifications is a human (e.g., a classifier in a library classification), whereas ontologies are primarily used by machines and, as such, they are the key enablers of the Semantic Web. Moreover,

designing a classification is part of everyday practice of many computer users, whereas designing a full-fledged ontology (expressed, for example, in OWL-DL) is a difficult and error-prone task even for ontology experts [17];

- **Purpose:** classifications are primarily used for the organization of (large) document collections into categories and subcategories such that these documents can be easily accessed by a human through browsing the classification tree in a top-down fashion. Ontologies are primarily used for modeling a particular domain such that the resulting model represents a shared view of a group of individuals [16];
- **Language:** as from the definition, classifications use natural language to describe nodes' categories. Natural language is well understood by humans but, due to its ambiguous nature, it is hard to be “understood” and reasoned about by machines. In contrast, ontologies are codified in a formal language which is unambiguously interpreted by machines. In fact, because ontologies are expressed in a formal language, they are often used for automated *reasoning* about the domain they model. Natural language is used in ontologies in a limited extent (e.g., to describe concept names) and, in general, has no functional value in reasoning operations on ontologies;
- **Nodes:** in an ontology, nodes normally represent atomic concepts (e.g., *car*, *wine*), whose names are shown next to the corresponding nodes when ontologies are visualized. In a classification, a label can represent a rather complex concept (e.g., “Open Source and Linux in Education”) or an individual (e.g., “Napoleon Bonaparte”);
- **Edges:** in an ontology graph, edges have a well defined semantics and they usually encode **sub-class-of**, **part-of** and other relations that hold between the two concepts connected with by an edge. In a classification, an edge implicitly represents either: (i) a *specification* relation which can be thought of as an **is-a** relation (e.g., an edge from “Animals” to “Humans”) or as a **part-of** relation (e.g., an edge from “Europe” to “Italy”); or, (ii) a *facet* relation which encodes the fact that the label of the child node represents an aspect of meaning of the parent node (e.g., an edge from “Animals” to “Images”) [3]. It is a bad practice to connect two nodes whose labels denote disjoint concepts (e.g., “non-living things” and “living organisms”) as in this case the child node and all its descendants cannot be populated with any document in a meaningful way;
- **Instances:** in an ontology, node instances are representatives of the node class and of all its ancestor classes in the **sub-class-of** hierarchy. They are in the **instance-of** relation with the class(es) they belong to. In a classification, node instances are not necessarily representatives of the class denoted by the node label, and can be documents which are about objects described by the set of labels of the nodes on the path from the given node to the root. For example, a node labeled “birds” may be populated with pictures of birds if the label of the parent node is “pictures”.

As shown above, classifications and ontologies are quite different and they have their cons and pros with respect to each other. We summarize their dis-

tinguishable features in Table 1 and, in the next section, we show how we can bridge the gap between them thus combining their pros within a single knowledge representation structure.

Table 1. Comparison between classification schemes and ontologies

Category	Classification Schemes	Ontologies
Users	Humans	Machines
Purpose	Organization of (large) document collections	Modeling of a domain
Language	Natural language, e.g. English	Formal language, e.g. OWL
Nodes	Usually represent complex concepts or individuals	Usually represent atomic concepts
Edges	Do not have well defined semantics	Have well defined semantics
Instances	Are not necessarily instances of the class in which they are populated	Are instances of the class in which they are populated
Examples	DDC, LCC, Colon classification	Gene ontology ^a , OpenCyc ontology ^b , MeSH ontology

^a <http://www.geneontology.org/>

^b <http://www.opencyc.org/>

3 From Classifications to OWL Ontologies

In this section we show how a classification, as defined in Section 2, can be converted into an OWL ontology. Particularly, we show how classification elements, namely: labels, nodes, edges, documents, and document-node links are encoded into OWL structures. Note that encoding classification labels requires converting from a natural language to a formal language, whereas encoding classification nodes and edges requires only structural manipulation. In Section 3.1 we discuss how we solve the former problem and in Section 3.2 we show how we solve the latter one. In Section 3.3 we show how we encode classification documents and document-node links as class instances. In Section 3.4 we show how the resulting OWL ontology can be enriched with a set of axioms such that it can be better suited for automated reasoning. Finally, in Section 3.5 we discuss which subset of the OWL language is required in order to encode classifications into ontologies.

3.1 From Labels to Concepts of Labels

In the conversion of natural language labels into a formal language we follow the approach presented in [4], which describes how these labels can be converted into a propositional concept language. The underlying idea of this approach is that senses of words, appearing in a label, are converted into atomic concepts, whereas punctuation and syntactic relations between words in the label are converted into

logical connectives (such as conjunction $\sqcap$ and disjunction $\sqcup$) and parenthesis. As discussed in [9], the extension of these concepts is the set of documents about the objects or individuals referred to by the (lexically defined) concepts. As shown in the same article, this interpretation has some nice properties such as it provides the possibility to represent individuals as concepts, and not as instances (e.g., the extension of concept `George.Bush` is the set of documents about the president George Bush), and to treat classification edges as the intersection of concepts. In the analysis of natural language labels we exploit the natural language processing (NLP) pipeline presented in [21]. Differently from the standard approaches to NLP, this pipeline is adapted to be applied on web directory labels. In the following we present the main steps of the pipeline and we show how we complete some of them with the conversion to OWL.

1. Sense retrieval. At this step, we retrieve the senses of each word in the label from the WordNet lexical database [15]. Apart from this, we identify words which are not found in WordNet.

2. Sense disambiguation. At this step, we leave only one sense per ambiguous word following the word sense disambiguation algorithm presented in [21]. The algorithm exploits the structure of the classification, WordNet relations such as hypernymy, and the most frequent sense heuristic to disambiguate the meaning.

3. Building atomic concepts. At this step we convert the disambiguated senses as well as the words which are not found in WordNet into atomic concepts and encode them as OWL classes. Following the approach described in [19], we define the URI scheme to uniquely identify OWL classes generated from WordNet senses as follows:

`Synset- + lexical_form_of_the_word- + POS- + synset_number`

where `synset_number` is the number of the synset⁷ to which the sense belongs in WordNet, and `lexical_form_of_the_word` is the lemma of the *first* word in the given synset. This allows us to represent synonymous words as one OWL class and not as multiple classes with equivalence relations defined between them.

For example, the URI for the atomic concept `java` which is generated from the sense *coffee* of the noun `java` is: `Synset-Coffee-Noun-41492`. An OWL class for this atomic concept is defined as follows:

`<owl:Class rdf:ID="Synset-Coffee-Noun-41492" />`

We form URIs for the words which are not found in WordNet as their literal representation in the label. For example, word *xyz* is encoded as an OWL class with URI `"xyz"`. This allows us to encode unknown words with the same spelling as one OWL class. Note that the encoding of words into concepts is done in a way to build a minimal set of concepts. The main reason for this choice is efficiency.

4. Building complex concepts. At this step we build complex concepts from atomic concepts following the approach discussed in [4]. For instance, a label composed of a sequence of adjectives followed by a noun group is converted into the logical conjunction ($\sqcap$) of the concepts corresponding to the adjectives

⁷ In WordNet, a *synset* is a set of one or more synonymous words which is assigned a unique numeric identifier, a gloss, and other metadata [15].

and to the nouns; prepositions like “of” and “in” are converted into the logical conjunction; coordinating conjunctions “and” and “or” are converted into the logical disjunction ($\sqcup$), and so on.

We convert complex concepts generated from the labels into classes in OWL. We define the following URI schema to uniquely identify these OWL classes:

Label- + node_label- + node_number

where **node_label** is the label of the node without spaces, and where each word starts with a capitalized letter. For example, the URI for the label “Society and Law Culture” of node 2 of the classification given in Figure 1 is: **Label-SocietyAndLawCulture-2**. The OWL class for this label is as follows:

```
<owl:Class rdf:ID="Label-SocietyAndLawCulture-2">
  <owl:unionOf rdf:parseType="Collection">
    <owl:Class rdf:ID="Synset-Society-Noun-318"/>
    <owl:intersectionOf rdf:parseType="Collection">
      <owl:Class rdf:ID="Synset-Law-Noun-51793"/>
      <owl:Class rdf:ID="Synset-Culture-Noun-38542"/>
    </owl:intersectionOf>
  </owl:unionOf>
</owl:Class>
```

3.2 From Concepts at Labels to Concepts at Nodes

As discussed in Section 2, edges in a classification represent either a specification or a facet relation, which can be generalized to the following observation: the meaning of a child node consists of what the meaning of its label and the meaning of the parent node have in common. We formalize this observation in the notion of *concept of node* [5, 8, 6, 7], which is defined as follows:

$$C_i = \begin{cases} l_i^F & \text{if } n_i \text{ is the root of } C \\ l_i^F \sqcap C_j & \text{if } n_i \text{ is not the root of } C, \text{ where } n_j \text{ is the parent of } n_i \end{cases} \quad (1)$$

where C_i is the concept of node n_i and l_i^F is the concept of label of node n_i . Concepts at nodes are converted into classes in OWL. The URI schema used to uniquely identify OWL classes corresponding to nodes is defined as follows:

Node- + node_label- + node_number

For example, in Figure 1 the URI for the root node labeled “Top” with id 1 is: **Node-Top-1**. An OWL class for this root node is built as follows:

```
<owl:Class rdf:ID="Node-Top-1">
  <equivalentClass rdf:resource="#Label-Top-1"/>
</owl:Class>
```

The URI for node 16 labeled “Programming Language” is: **Node-ProgrammingLanguage-16** and its corresponding OWL class is built as follows:

```
<owl:Class rdf:ID="Node-ProgrammingLanguage-16">
```

```

<owl:intersectionOf rdf:parseType="Collection">
  <owl:Class rdf:ID="Label-ProgrammingLanguage-16"/>
  <owl:Class rdf:ID="Node-Computer-15"/>
</owl:intersectionOf>
</owl:Class>

```

Note that classification edges are implicitly encoded in the definitions of OWL classes representing concepts at nodes. Namely, since these classes are defined as the intersection of the concept at node of the parent and the concept at label of the child node, then the structure of the classification can be reconstructed by analyzing node class definitions.

3.3 From Documents to Class Instances

We convert a document into an instance of the OWL Thing class. We assume that each document has a URL and we use it to uniquely identify the corresponding instance in OWL. Moreover, if a document has a title and a description (as web directory documents normally have), then we encode them in `rdfs:label` and `rdfs:comment` properties accordingly. For example, a document with URL <http://java-source.net/>, with title “Java Open Source Software”, and with description “A directory of open source software focused on Java” is encoded in OWL as follows:

```

<owl:Thing rdf:about="#http://java-source.net/">
  <rdfs:label>Java Open Source Software</rdfs:label>
  <rdfs:comment>A directory of open source software focused on Java
</rdfs:comment>
</owl:Thing>

```

We convert document-node links of a document by defining the `rdf:type` relation from the instance, representing the document, to the class(es) representing the node(s) in which the document is classified. For instance, if the above mentioned document is classified in nodes 2 and 4 of the classification shown in Figure 1, then these document-node links are encoded as follows:

```

<owl:Thing rdf:about="#http://www.laweasy.com">
  <rdf:type rdf:resource="#Node-SocietyAndLawCulture-2"/>
  <rdf:type rdf:resource="#Node-Law-4"/>
</owl:Thing>

```

3.4 Semantic Enrichment

Since OWL classes, which correspond to word senses, are mapped to synsets in WordNet, we can exploit the relations between synsets and relations between words within synsets in order to enrich the resulting OWL ontologies with additional relations between classes. The enrichment is based on these two rules:

- **Rule 1:** In WordNet, synsets are organized into hierarchies based, for example, on the hypernym (i.e., *is-a* or *is-kind-of*) relation [15]. For instance, the

synset denoting “Java” (as “a simple platform-independent object-oriented programming language”) has a hypernym synset denoting “object-oriented programming language” (as “a programming language that enables the programmer to associate a set of procedures with each type of data structure”). If two OWL classes (*c1-1* and *c1-2*) correspond to two senses (*sen-1* and *sen-2*) belonging to two synsets (*syn-1* and *syn-2*) among which there is a hypernym relation defined in WordNet (e.g., *syn-2* is a hypernym for *syn-1*), then we define an `rdfs:subClassOf` relation between these two classes (i.e., *c1-1* `rdfs:subClassOf` *c1-2*) as follows:

```
<owl:Class rdf:about="#Synset-Java-Noun-41493">
  <rdfs:subClassOf rdf:resource="#Synset-ProgrammingLanguage-Noun-45-219"/>
</owl:Class>
```

- **Rule 2:** Antonym relations in WordNet are defined among *words* within synsets (and not among synsets). We translate these relations into `owl:disjointWith` relations among classes corresponding to senses of the two antonym words. For instance, the antonym of the word “day” in the synset {day, daytime, daylight} is the word “night” in the synset {night, nighttime, dark}. The former synset is the third sense of the noun “day” and the latter synset is the first sense of the noun “night”. Classes, associated with these two senses, are declared to be disjoint as follows:

```
<owl:Class rdf:about="#Synset-Day-Noun-12826">
  <owl:disjointWith rdf:resource="#Synset-Night-Noun-38819"/>
</owl:Class>
```

The enrichment of classification OWL ontologies according to the two rules described above allows us to make these ontologies more suitable for reasoning as the underline axiom base grows.

3.5 OWL Sublanguage

OWL ontologies, generated from classifications, fall into the OWL Lite or OWL DL subset of OWL. There are two factors which require OWL DL:

- the logical disjunction that may appear after the conversion of natural language labels and which is converted into the `owl:unionOf` construct;
- disjoint axioms that may appear at the semantic enrichment step and which are converted into the `owl:disjointWith` construct.

Both above mentioned constructs are forbidden in OWL Lite. Note that the conversion to OWL does not require the use of constructs of OWL Full which leaves us within a decidable subset of OWL.

4 Evaluation

To evaluate our approach, we selected four subtrees with the maximum depth of 3 from the DMOz web directory. In Table 2 we report statistical data of the datasets. There are 476 nodes in the selected subtrees, which have 548 tokens in total, out of which, 527 tokens are found in WordNet (i.e., WordNet coverage is 96.17%). Out of the set of words found in WordNet, 223 (i.e., 42.31%) are ambiguous with the average polysemy of 3.36. In our experiments we used WordNet version 2.0.

Table 2. Statistics of the dataset

Dataset	Nodes	Average Branching Factor	Average Subtree Depth	Tokens Per Label	Words with Senses in WordNet	Noun Senses	Adjective Senses
Countries ^a	245	6.26	3	1.07	261	256	5
Europe ^b	75	4.22	3	1.12	86	86	0
Asia ^c	76	4.24	3	1.18	89	88	1
Africa ^d	80	4.31	3	1.15	94	93	1

^a <http://dmoz.org/Regional/Countries/>.

^b <http://dmoz.org/Regional/Europe/>.

^c <http://dmoz.org/Regional/Asia/>.

^d <http://dmoz.org/Regional/Africa/>.

4.1 Correctness

We evaluated the most critical step of the NLP pipeline, i.e., the word sense disambiguation (see Section 3.1) algorithm, whose performance results are reported in Table 3. The accuracy of this step largely affects the correctness of the results of reasoning on these OWL ontologies, as we show in Section 5.5.

Table 3. Accuracy of the word sense disambiguation algorithm

Dataset	Ambiguous Tokens	Disambiguation Accuracy(%)
Countries	92	76.54
Europe	38	77.01
Asia	47	80.89
Africa	46	79.13

4.2 OWL Sublanguage

In Table 4 we report statistical data for the generated OWL ontologies and in Table 5 we provide details on the kind and number of axioms before and after semantic enrichment.

Table 4. Statistics of the generated OWL ontologies

Ontology	Nodes	Sense Classes	Label Classes	Node Classes	Class Axioms	Individual Axioms	intersectionOf Constructs	unionOf Con- structs
Countries	245	261	245	245	873	0	265	4
Europe	75	86	75	75	155	183	76	10
Asia	76	89	76	76	203	125	80	9
Africa	80	94	80	80	212	253	84	9

Table 5. Axioms before and after semantic enrichment

Ontology	Equivalent Class Axioms		SubClass Axioms		Disjoint Class Axioms		Individual Axioms	
	Before	After	Before	After	Before	After	Before	After
Countries	490	490	0	383	0	0	0	0
Europe	152	152	0	3	0	0	183	183
Asia	152	152	0	51	0	0	125	125
Africa	160	160	0	52	0	0	253	253

Noteworthy, most of the constructs in the generated ontologies are valid in OWL Lite. There are only few `owl:unionOf` constructs, which require the use of OWL DL for the representation of these ontologies.

5 Optimizing Classifications

In this section we show some practical examples of reasoning on classification OWL ontologies. For instance, we show how they can be checked for consistency, how their structure can be rationalized, and how nodes with similar contents to a given node can be found.

5.1 Consistency

We used Protégé OWL Plugin [14] and its reasoning capabilities to detect logical inconsistencies within the classification OWL ontologies. We used reasoning capabilities of both Pellet 1.5 and Fact++ OWL reasoners launched with Protégé.

None of the reasoners reported that the classification OWL ontologies were inconsistent.

5.2 Rational Forms

Classifications may not be perfect. For this reason we may need to reconstruct a classification based on the “most specific subsumer” relation. Nodes get parents which most specifically describe them, still being more general. The new structure is called, a *rational form* of a classification. The idea behind the rationalization of classifications is to build a classification which better corresponds to a taxonomic structure. At the semantic enrichment step rational form is built. The classification given in Figure 3(b) is a rational form of the classification given in Figure 3(a). This rational form is obtained as “programming language” is a direct hypernym of “object-oriented programming language” in WordNet and that is computed and then converted to `owl:subClassOf` relation between them at semantic enrichment step. Note that classification semantics does not change when going from classification to rational form of classification as the set of concepts at nodes remains the same.




Fig. 3. (a) Classification; (b) Rational form of the classification given in (a)

5.3 Minimizing Effort

In Table 6 we report the kind and number of found relations within and across the four ontologies. For example, the reasoner found an equivalent relation between node class */Regional/Countries/Italy* and node class */Regional/Europe/Italy*. This is an example of how reasoning on classification OWL ontologies can help web directory editors find interrelated parts of the web directory and, thus, improve its organizational structure without manual inspection. Note that no disjointness relations were found because we did not have disjoint axioms in the produced OWL ontologies.

Table 6. Found relations within and across the four ontologies

Ontology	Countries		Europe		Asia		Africa	
	$\sqsubseteq$	$\equiv$	$\sqsubseteq$	$\equiv$	$\sqsubseteq$	$\equiv$	$\sqsubseteq$	$\equiv$
Countries	383	490	386	642	392	642	387	650
Europe	386	642	3	152	51	304	52	312
Asia	392	642	51	304	51	152	100	312
Africa	387	650	52	312	100	312	52	160

5.4 Computing See-Also Links

Apart from the four ontologies, we experimented with another classification OWL ontology and we observed that the individuals asserted to the OWL class which corresponds to the classification node */Games and Activities/Kids and Teens/Football* are inferred as the individuals of the OWL class which corresponds to the classification node */Sports Athletics Funs/Youth and High School/Soccer*, and vice versa. This kind of reasoning can be used for finding similar documents populated in different nodes, which will help in building *see-also* links.

5.5 Errors

Apart from correct relations, we found also some incorrect ones. For example, the reasoner found an erroneous more specific relation between node class */Regional/Europe/Georgia* and node class */Regional/Countries/United States*. As discussed earlier, this problem is caused by the lack of accuracy of the word sense disambiguation algorithm. Evaluating the correctness and completeness characteristics of the computed set of relations between ontology classes is outside the scope of the current paper. Interested readers are referred to [5] for a complete account.

6 Related Work

The current work is a representative of a recent trend in the Semantic Web community towards the use of *lightweight semantics* (as opposed to expressive logic languages) and *lightweight ontologies* [9] (as opposed to full-fledged ontologies), the generation of which can be potentially supported by ordinary users which constitute the long tail of the Semantic Web. The trend has been formed through a number of scientific publications (e.g., see [20, 4, 18, 13]) and is currently supported by a number of R&D projects (e.g., MATURE⁸, OpenKnowledge⁹) and systems (e.g., OntoWiki¹⁰). The current work contributes to this trend by proposing an approach in which classifications, which are often called

⁸ MATURE, Integrated Project (IP), FP7-216356, see <http://mature-ip.eu>.

⁹ OpenKnowledge, STREP, FP6-27253, see <http://www.openk.org/>

¹⁰ OntoWiki, see <http://ontowiki.net/Projects/OntoWiki>.

(informal) lightweight ontologies [9] and whose most representative instantiations on the web are web directories, can be automatically converted into formal OWL ontologies, ready to be embedded in Semantic Web applications.

There are few lines of work which are close in spirit to our approach. For instance, in [20], the authors propose a method to converting thesauri to OWL ontologies in which they provide a detailed account of how elements of a thesaurus are converted into OWL structures. This approach is based on a manual analysis of thesauri, whereas our approach allows for a fully automatic conversion. Another approach, discussed in [18], comes from the Digital Library community and presents a conceptual structure and transition procedure to support the shift from a traditional knowledge organization system (KOS) and, particularly, a thesaurus, towards a full-fledged and semantically rich KOS. While providing an in-depth analysis of the shortcomings of the traditional KOSs and of the benefits of semantic KOSs as well as providing a set of rules for converting thesaurus elements into ontology constructs, the approach lacks a specification of how a KOS can be converted into an ontology language, such as OWL – the ultimate conversion step discussed in detail in the current paper.

The approach described in [12] allows us to convert a hierarchical classification into an OWL ontology by deriving OWL classes from classification labels and by arranging these classes into a hierarchy (based on the `rdfs:subClassOf` relation) following the classification structure. The approach is based on some application-dependent assumptions such as that one label represents one atomic concept, and that relations between labels can be defined as `sub-class-of` relations in some particular context (e.g., concept “ice” is more specific than concept “non-alcoholic beverages” when considered in the context of procurement). These assumptions do not hold in a general case and are not made in our approach. Apart from this, our approach differs from [20, 18, 12] in that it is generic and, therefore, suitable for automatic conversion in OWL of any knowledge representation structure whose core can be represented in the form of a classification as defined in this paper.

7 Conclusions

In this paper we have presented a fully automated approach to converting generic classification schemes into OWL ontologies. The proposed approach allows us to leverage on top of classifications, being the interfaces to knowledge for humans, and ontologies, being the interfaces to knowledge for machines on the Semantic Web. Furthermore, as shown above, our approach provides immediate advantage and it allows to help the user in building better classifications more suited for reasoning. Potentially, the approach allows for a cost-free seamless integration of a vast amount of classification structures on the web and in personal repositories into the Semantic Web infrastructure, thus reducing the problem of the lack of semantically rich data. The first experimental results, reported in this paper, show that reasoning on classification OWL ontologies can be used for building practical Semantic Web applications.

References

1. T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, (284(5)):34–43, May 2001.
2. S. Bechhofer et al. OWL Web ontology language reference, W3C recommendation, February 2004.
3. F. Giunchiglia, M. Marchese, and I. Zaihrayeu. Towards a theory of formal classification. In *Proceedings of the AAAI-05 Workshop on Contexts and Ontologies: Theory, Practice and Applications (C&O-2005)*, Pittsburgh, Pennsylvania, USA, 2005.
4. F. Giunchiglia, M. Marchese, and I. Zaihrayeu. Encoding classifications into lightweight ontologies. In *Journal on Data Semantics (JoDS) VIII*, Winter 2006.
5. F. Giunchiglia and P. Shvaiko. Semantic matching. *Knowledge Engineering Review*, 18(3):265–280, 2003.
6. F. Giunchiglia, P. Shvaiko, and M. Yatskevich. Semantic schema matching. In *OTM Conferences (1)*, pages 347–365, 2005.
7. F. Giunchiglia and M. Yatskevich. Element level semantic matching. In *Meaning Coordination and Negotiation workshop, ISWC*, 2004.
8. F. Giunchiglia, M. Yatskevich, and E. Giunchiglia. Efficient semantic matching. In *ESWC*, pages 272–289, 2005.
9. F. Giunchiglia and I. Zaihrayeu. Lightweight ontologies. In *The Encyclopedia of Database Systems, to appear*. Springer, 2008.
10. T. R. Gruber. A translation approach to portable ontology specifications. *Knowl. Acquis.*, 5(2):199–220, 1993.
11. N. Guarino. Some ontological principles for designing upper level lexical resources. In *First International Conference on Lexical Resources and Evaluation*, volume 2830, Granada, Spain, May 1998.
12. M. Hepp. Representing the hierarchy of industrial taxonomies in OWL: The gen/tax approach. In *Proceedings ISWC Workshop on Semantic Web Case Studies and Best Practices for eBusiness (SWCASE05)*, 2005.
13. M. Hepp and J. de Bruijn. Gen tax: A generic methodology for deriving OWL and RDF-S ontologies from hierarchical classifications, thesauri, and inconsistent taxonomies. In *ESWC*, 2007.
14. H. Knublauch, M. A. Musen, and A. L. Rector. Editing description logic ontologies with the Protégé OWL plugin. In *Description Logics*, 2004.
15. G. Miller. *WordNet: An electronic Lexical Database*. MIT Press, 1998.
16. H. S. Pinto, S. Staab, and C. Tempich. Diligent: Towards a fine-grained methodology for distributed, loosely-controlled and evolving engineering of ontologies. In *ECAI*, pages 393–397, 2004.
17. A. L. Rector, N. Drummond, M. Horridge, J. Rogers, H. Knublauch, R. Stevens, H. Wang, and C. Wroe. OWL pizzas: Practical experience of teaching OWL-DL: Common errors & common patterns. In *EKAU*, pages 63–81, 2004.
18. Dagobert Soergel, Boris Lauser, Anita Liang, Frehiwot Fisseha, Johannes Keizer, and Stephen Katz. Reengineering thesauri for new applications: The agrovoc example. *J. Digit. Inf.*, 4(4), 2004.
19. M. van Assem, A. Gangemi, and G. Schreiber. RDF/OWL representation of WordNet, W3C working draft, June 2006.
20. M. van Assem, M. R. Menken, G. Schreiber, J. Wielemaker, and B. J. Wielinga. A method for converting thesauri to RDF/OWL. In *International Semantic Web Conference*, pages 17–31, 2004.

21. I. Zaihrayeu, L. Sun, F. Giunchiglia, W. Pan, Q. Ju, M. Chi, and X. Huang. From web directories to ontologies: Natural language processing challenges. In *ISWC/ASWC*, pages 623–636, 2007.