

# Formalizing the Get-Specific Document Classification Algorithm

Fausto Giunchiglia, Ilya Zaihrayeu, and Uladzimir Kharkevich

Department of Information and Communication Technology  
University of Trento, Italy  
{fausto,ilya,kharkevi}@dit.unitn.it

**Abstract.** The paper represents a first attempt to formalize the get-specific document classification algorithm and to fully automate it through reasoning in a propositional concept language without requiring user involvement or a training dataset. We follow a knowledge-centric approach and convert a natural language hierarchical classification into a formal classification, where the labels are defined in the concept language. This allows us to encode the get-specific algorithm as a problem in the concept language. The reported experimental results provide evidence of practical applicability of the proposed approach.

## 1 Introduction

Classification hierarchies have always been a natural and effective way for humans to organize their knowledge about the world. These hierarchies are rooted trees where each node defines a topic category. Child nodes' categories define aspects or facets of the parent node's category, thus creating a multifaceted description of the objects which can be classified in these categories. Classification hierarchies are used pervasively: in conventional libraries (e.g., the Dewey Decimal Classification system (DDC) [8]), in web directories (e.g., DMOZ [2]), in e-commerce standardized catalogues (e.g., UNSPSC [3]), and so on.

Standard classification methodologies amount to manually organizing objects into classification categories following a predefined system of rules. The rules may differ widely in different approaches, but there is one classification pattern which is commonly followed. The pattern is called the *get-specific principle*, and it requires that an object is classified in a category (or in a set of categories), which most specifically describes the object. Following this principle is not easy and is constrained by a number of limitations, discussed below:

- the meaning of a given category is implicitly codified in a natural language label, which may be ambiguous and may therefore be interpreted differently by different classifiers;
- a link, connecting two nodes, may also be ambiguous in the sense that it may specify the meaning of the child node, of the parent node, or of both;
- as a consequence of the previous two items, the classification task also becomes ambiguous in the sense that different classifiers may classify the same objects differently, based on their subjective opinion.

In the present paper we propose an approach to converting classifications into *formal classifications*, whose labels are encoded in a propositional concept language. Apart from this, we present a classification model and show how the get-specific algorithm can be described in this model. We then show how the model and the algorithm can be encoded in the concept language, which allows us to fully automate document population in formal classifications through propositional reasoning. Note that by doing this, we eliminate the three ambiguities discussed above. In order to evaluate our approach, we have re-classified documents from several branches of the DMOz directory without any human involvement or an a priori created training dataset. The results show the viability of the proposed approach, which makes it a realistic alternative to the standard classification approaches used in Information Science.

The remainder of the paper is organized as follows. In Section 2 we introduce the classification model, we show how the get-specific algorithm can be described in this model, and we identify the main problems peculiar to the algorithm. In Section 3 we show how classifications can be translated into formal classifications, how the get-specific algorithm can be encoded in the concept language, and how its peculiar problems can be dealt with in the concept language. In Section 4 we present and discuss evaluation results of our approach. In Section 5 we discuss the related work and, in particular, we compare our approach to that used in Information Science. Section 6 summarizes the results and concludes the paper.

## 2 The Get-Specific Classification Algorithm

Classifications are hierarchical structures used for positioning objects in such a way, that a person, who navigates the classifications, will be facilitated in finding objects related to a given topic. To attain such organization of objects, in standard classification approaches, objects are manually classified by human classifiers which follow a predefined system of rules. The actual system of rules may differ widely in different classification approaches, but there are some generic principles which are commonly followed. These principles make the ground of the *get-specific* algorithm, described in the rest of this section.

### 2.1 Classifications and a Classification Model

To avoid ambiguity of interpretation, in Definition 1 we formally define the notion of classification; and in Figure 1 we give an example of a classification, extracted from the DMOz web directory and adjusted for sake of presentation.

**Definition 1.** *A classification is a rooted tree  $C = \langle N, E, L \rangle$  where  $N$  is a set of nodes,  $E$  is a set of edges on  $N$ , and  $L$  is a set of labels expressed in a natural language, such that for any node  $n_i \in N$ , there is one and only one label  $l_i \in L$ .*

We see the process of classification as a decision making procedure in which the classification tree is divided into a set of minimal decision making blocks. Each block consists of a node (called the root node of the block) and its child




Fig. 1. A part of the DMoz web directory

nodes (see Figure 2). While classifying an object, the classifier considers these blocks in a top-down fashion, starting from the block at the classification root node and then continuing to blocks rooted at those child nodes, which were selected for *further consideration*. These nodes are selected following decisions which are made at each block along two dimensions: *vertical* and *horizontal*. In the vertical dimension, the classifier decides which of the child nodes are selected as candidates for further consideration. In the horizontal dimension, the classifier decides which of the candidates are *actually* selected for further consideration. If none of the child nodes are appropriate or if there are no child nodes, then the root node of the block becomes a *classification alternative* for the given object. The process reiterates and continues recursively until no more nodes are left for further consideration. At this point, all the classification alternatives are computed. The classifier then decides which of them are most appropriate for the classification of the object and makes *the final classification choice*.




Fig. 2. The decision making block

## 2.2 Modelling the Get-Specific Classification Algorithm

In this subsection we discuss the general principles lying behind the get-specific algorithm and we show how these principles can be implemented within the model introduced in the previous subsection. Particularly, we discuss how vertical and horizontal choices, as well as the final classification choice are made.

- **Vertical choice.** Classification hierarchies are organized such that upper level categories represent more general concepts, whereas lower level categories represent more specific concepts. When the classifier searches for an

appropriate category for the classification of an object, she looks for the ones which most specifically describe the object and, therefore, when making a vertical choice, she selects a child node as a candidate if it describes the object more specifically than the parent does. For example, if a document about ancient Greek coins is classified in the classification from Figure 1, then node  $n_6$  is more appropriate for the classification than node  $n_5$ . When this principle is applied recursively, it leads to the selection of the category which lies as deep in the classification hierarchy as possible. The principle described above is commonly called the *get-specific principle*. Let us consider, for instance, how Yahoo! describes it:

*“When you suggest your site, get as specific as possible. Dig deep into the directory, looking for the appropriate sub-category.” [4]*

- **Horizontal choice.** Child nodes may describe different aspects or *facets* of the parent node and, therefore, more than one child node may be selected in the vertical choice if a multifaceted document is being classified. As a consequence of this, the classifier needs to decide which of the several sibling nodes are appropriate for further consideration. When one sibling node represents a more specific concept than another, then the former is usually preferred over the latter. For example, node  $n_{10}$  is more appropriate for the classification of ancient Greek coins than node  $n_{12}$ . As a rule of thumb, the horizontal choice is made in favor of as few nodes as possible and, preferably, in favor of one node only. We call the principle described above, the *get-minimal principle*. Consider, for instance, how DMoz describes it.

*“Most sites will fit perfectly into one category. ODP categories are specialized enough so that in most cases you should not list a site more than once.” [1]*

- **Tradeoff between vertical and horizontal choices.** The two principles described above cannot always be fulfilled at the same time. Namely, if the vertical choice results in too many candidates, then it becomes hard to fulfill the principle of minimality in the horizontal choice. In order to address this problem, a tradeoff needs to be introduced between the two requirements, which usually means trading specificity in favor of minimality. The following is an example of a tradeoff rule used in DMoz:

*“If a site offers many different things, it should be listed in a more general category as opposed to listing it in many specialized subcategories.” [1]*

- **The final classification choice.** When all classification alternatives are determined, the classifier confronts all of them in order to make her final classification choice. Note that now the choice is made not at the level of a minimal decision making block, but at the level of the whole classification. However, the classifier uses the same selection criteria as those used in the horizontal choice. For example, nodes  $n_6$  and  $n_{13}$  are more appropriate for the classification of documents about ancient Greek coins than node  $n_{11}$ .

### 2.3 Problems of the Get-Specific Classification Algorithm

As discussed in [10], there are several problems which are common to document classification algorithms. The problems are caused by the potentially large size of classifications, by ambiguity in natural language labels and in document descriptions, by different interpretations of the meaning of parent-child links, and so on. All these problems lead to nonuniform, duplicate, and error-prone classification. In addition to the problems discussed in [10], the get-specific algorithm has two peculiar problems, related to the two decision dimensions. We discuss these problems below on the example of a document titled “*Gold Staters in the Numismatic Marketplace*”, being classified in the classification from Figure 1.

- **Vertical choice: the “I don’t know” problem.** The classifier may make a mistake because she does not (fully) understand the meaning of a child node or the relation of the document to that node, whereas the node is a valid candidate. For example, the classifier may not know that “Gold Stater” is a coin of ancient Greece and, therefore, will erroneously classify the document into node  $n_5$ , whereas a more appropriate node is  $n_6$ .
- **Horizontal choice: the “Polarity change” problem.** The classifier may make a mistake when one of the sibling candidate nodes is more appropriate for further consideration than another, but a descendent of the latter is more appropriate for the classification than a descendant of the former node. For instance, the label of node  $n_{10}$  more specifically describes the document than the label of node  $n_{12}$ . Therefore, the classifier will choose node  $n_{10}$  only as a candidate and will finally classify the document in node  $n_{11}$ , whereas a more appropriate node for the classification is node  $n_{13}$ , a descendent of  $n_{12}$ .

## 3 Formalizing the Get-Specific Classification Algorithm

In this section we formalize the get-specific classification algorithm by encoding it as a problem expressed in propositional Description Logic language [5], referred to as  $L^C$ . First, we discuss how natural language node labels and document descriptions are converted into formulas in  $L^C$ . Second, we discuss how we reduce the problems of vertical, horizontal, and final classification choices to fully automated propositional reasoning. Finally, we show how the problems discussed in Section 2.3 can be dealt with in a formal way.

### 3.1 From Natural Language to Formal Language

Classification labels are expressed in a natural language, which is ambiguous and, therefore, is very hard to reason about. In order to address this problem, we encode classification labels into formulas in  $L^C$  following the approach proposed in [10]. This allows us to convert the classification into a new structure, which we call *Formal Classification* (FC):

**Definition 2.** A *Formal Classification* is a rooted tree  $FC = \langle N, E, L^F \rangle$  where  $N$  is a set of nodes,  $E$  is a set of edges on  $N$ , and  $L^F$  is a set of labels expressed in  $L^C$ , such that for any node  $n_i \in N$ , there is one and only one label  $l_i^F \in L^F$ .

Note that even if  $L^C$  is propositional in nature, it has a set-theoretic semantics. As proposed in [10], the interpretation of a concept is the set of documents, which are *about* this concept. For instance, the interpretation of concept **Capital** (defined as “a seat of government”) is the set of documents about capitals, and *not* the set of capitals which exist in the world.

Below we briefly describe how we convert natural language labels into formulas in  $L^C$ . Interested readers are referred to [10] for a complete account. Figure 3 shows the result of conversion of the classification from Figure 1 into a FC.

1. **Build atomic concepts.** Senses of nouns and adjectives become atomic concepts, whose interpretation is the set of documents about the entities or individual objects, denoted by the nouns, or which possess the qualities, denoted by the adjectives. We enumerate word senses using WordNet [17], and we refer to them as follows: `pos-lemma#i`, where `pos` is the part of speech, `lemma` is the word lemma, and `i` is the sense number in WordNet.
2. **Disambiguate word senses.** Irrelevant word senses are identified and corresponding to them atomic concepts are discarded. As proposed in [15], if there exists a relation (e.g., synonymy, hypernymy, or holonymy) in WordNet between any two senses of two words in a label (or in different labels on a path to the root), then corresponding to them concepts are retained and other concepts are discarded. If no relation is found, then we check if a relation exists between two WordNet senses by comparing their glosses [13].
3. **Build complex concepts.** Complex concepts are built as follows: first, words’ formulas are built as the logical disjunction ( $\sqcup$ ) of atomic concepts corresponding to their senses (remaining after step 2). Second, syntactic relations between words are translated into logical connectives of  $L^C$ . For example, a set of adjectives followed by a noun group is translated into the logical conjunction ( $\sqcap$ ) of the formulas corresponding to the adjectives and to the nouns; prepositions like “of” and “in” are translated into the conjunction; coordinating conjunctions “and” and “or” are translated into the logical disjunction ( $\sqcup$ ); words and phrases denoting exclusions, such as “except” and “but not”, are translated into the logical negation ( $\neg$ ).

Before a document can be automatically classified, it has to be assigned an expression in  $L^C$ , which we call the document concept, written  $C^d$ . The assignment of a concept to a document is done in two steps: first, a set of  $n$  keyphrases is retrieved from the document using text mining techniques (see, for example, [22]); the keyphrases are converted to formulas in  $L^C$ , and the document concept is then computed as the conjunction of the formulas.

### 3.2 The Algorithm

In the following we describe how we make vertical and horizontal choices, compute the tradeoff, and make the final classification choice in FCs.




Fig. 3. Formal Classification

- **Vertical choice.** A child node  $n_i$  is a candidate, given that a document with concept  $C^d$  is being classified, if the label of the node,  $l_i^F$ , subsumes  $C^d$ , i.e., if the following holds:  $C^d \sqsubseteq l_i^F$ . In formulas, if  $N_c$  is the set of child nodes in the block, then we compute the vertical choice  $V(C^d)$  as:

$$V(C^d) = \{n_i \in N_c \mid C^d \sqsubseteq l_i^F\} \quad (1)$$

If the vertical choice results in no candidates, then root node  $n_r$  of the current block is added to the set of classification alternatives  $A(C^d)$ :

$$\text{if } |V(C^d)| = 0 \text{ then } A(C^d) \leftarrow A(C^d) \cup \{n_r\} \quad (2)$$

In Figure 4a we show an example of a situation when two child nodes  $n_2$  and  $n_4$  are selected for further consideration, and in Figure 4b we show an example of a situation when no child node can be selected.




Fig. 4. Vertical choice (“?” means no relation is found)

- **Horizontal choice.** Given the set of candidates  $V(C^d)$ , we exclude those nodes from the set, whose label is more general than the label of another node in the set. In formulas, we compute the horizontal choice  $H(C^d)$  as:

$$H(C^d) = \{n_i \in V(C^d) \mid \nexists n_j \in V(C^d), \text{ s.t. } j \neq i, l_j^F \sqsubseteq l_i^F, \text{ and } l_j^F \not\supseteq l_i^F\} \quad (3)$$

We introduce the last condition (i.e.,  $l_j^F \not\supseteq l_i^F$ ) to avoid mutual exclusion of nodes, whose labels in the FC are equivalent concepts. For instance, two syntactically different labels “seacoast” and “seashore” are translated into two equivalent concepts. When such situation arises, all the nodes, whose labels are equivalent, are retained in  $H(C^d)$ .

- **The tradeoff.** Whenever the size of  $H(C^d)$  exceeds some threshold  $k$ , the nodes of  $H(C^d)$  are discarded as candidates and root node  $n_r$  of the block is added to the set of classification alternatives  $A(C^d)$ . In formulas:

$$\mathbf{if} |H(C^d)| > k \mathbf{ then } H(C^d) \leftarrow \emptyset \mathbf{ and } A(C^d) \leftarrow A(C^d) \cup \{n_r\} \quad (4)$$

- **The final classification choice.** When no more nodes are left for further consideration, set  $A(C^d)$  includes all the classification alternatives. We compare them to make the final classification choice, but, differently from vertical and horizontal choices, we compare the *meanings* of nodes given their path to the root, and not their labels. We encode the meaning of node  $n_i$  into a concept in  $L^C$ , called *concept of node* [11], written  $C_i$ , and computed as:

$$C_i = \begin{cases} l_i^F & \text{if } n_i \text{ is the root of the } FC \\ l_i^F \sqcap C_j & \text{if } n_i \text{ is not the root, where } n_j \text{ is the parent of } n_i \end{cases} \quad (5)$$

Similar to how the horizontal choice is made, we exclude those nodes from  $A(C^d)$ , whose concept is more general than the concept of another node in the set. In formulas, we compute the final classification choice  $C(A)$  as:

$$C(A) = \{n_i \in A(C^d) \mid \nexists n_j \in A(C^d), \text{ s.t. } j \neq i, C_j \sqsubseteq C_i, \text{ and } C_j \not\sqsupseteq C_i\} \quad (6)$$

The last condition (i.e.,  $C_j \not\sqsupseteq C_i$ ) is introduced to avoid mutual exclusion of nodes with the same meaning in the classification hierarchy. For instance, two paths *top/computers/games/soccer* and *top/sport/soccer/computer\_games* lead to two semantically equivalent concepts. When such situation arises, all the nodes with the same meaning are retained in  $C(A)$ .

Computing Equations 1, 3, and 6 requires verifying whether the subsumption relation holds between two formulas in  $L^C$ . As shown in [10], a problem expressed in  $L^C$  can be rewritten as an equivalent problem expressed in propositional logic. Namely, if we need to check whether a certain relation *rel* (which can be  $\sqsubseteq$ ,  $\sqsupseteq$ ,  $\equiv$ , or  $\perp$ ) holds between two concepts  $A$  and  $B$ , given some knowledge base  $\mathcal{KB}$  (which represents our a priori knowledge), we construct a propositional formula according to the pattern shown in Equation 7 and check it for validity:

$$\mathcal{KB} \rightarrow \text{rel}(A, B) \quad (7)$$

### 3.3 Dealing with Problems

Encoding a classification algorithm into a problem in  $L^C$  allows it to avoid many problems, which are common to classification algorithms [10]. Particularly, since the problem is encoded in a formal language, there is no ambiguity in interpretation of classification labels, of edges, and document contents. Apart from this, since computation is performed by a machine, the problem of classification size becomes largely irrelevant. Finally, since the formal algorithm is deterministic, the classification is always performed in a uniform way.

In Section 2.3 we discussed two problems, peculiar to the get-specific algorithm. Below we discuss what they mean in  $L^C$  and how they can be dealt with.



- **Vertical choice: The “I don’t know” problem.** This problem arises when the specificity relation in Equation 1 cannot be computed while a human observes that it exists. The problem is caused by lack of background knowledge and it can be dealt with by adding missing axioms to the underlying knowledge base [12]. For instance, if we add a missing axiom which states that concept **Stater** (defined as “any of the various silver or gold coins of ancient Greece”) is more specific than concept **Greek** (defined as “of or relating to or characteristic of Greece . . .”), then the algorithm will correctly classify document “*Gold Staters in the Numismatic Marketplace*” into node  $n_6$  in the classification shown in Figure 1.
- **Horizontal choice: The “Polarity change” problem.** The problem arises when the label of node  $n_i$  is more specific than the label of its sibling node  $n_j$  (i.e.,  $l_i^F \sqsubseteq l_j^F$ ), but the concept of a  $n_i$ ’s descendant node  $n_k$  is more general than the concept of a  $n_j$ ’s descendant node  $n_m$  (i.e.,  $C_k \supseteq C_m$ ). In the simplest case, this problem can be dealt with by not performing the horizontal choice. In this case, both  $n_k$  and  $n_m$  will be in the classification alternative set for some document, and  $n_k$  will then be discarded when the final classification choice is made.

## 4 Evaluation

In order to evaluate our approach, we selected four subtrees from the DMoz web directory, converted them to FCs, extracted concepts from the populated documents, and automatically (re)classified the documents into the FCs. We extracted document concepts by computing the conjunction of the formulas corresponding to the first 10 most frequent words appearing in the documents (excluding stop words). We used WordNet 2.0 [17] for finding word senses and their relations, and we used S-Match [11] for computing Equation 7. Parameter  $k$  for tradeoff computation was set to 2.

In the evaluation we employ standard information retrieval measures such as micro- and macro-averaged precision, recall, and F1 [19]. In Table 1 we report dataset statistics and evaluation results for each of the four datasets. We performed a detailed analysis of the “Languages” dataset results (see Figure 5). In Figure 5a we show how precision and recall are distributed among nodes. Figure 5b shows how far (in terms of the number of edges) an automatically classified document is from the node where it was actually classified in DMoz.

From Figure 5a we observe that about 40% of nodes in the “Languages” dataset have precision and recall equal to 0<sup>1</sup>. After manual inspection of the results, we concluded that this problem is caused by lack of background knowledge. For instance, 8 documents about Slovenian language were misclassified because there was no WordNet synset “Slovenian” defined as “the Slavic language spoken in Slovenia” and a hypernym relation of it with synset “Slavic language”. Figure 5b shows that about 20% of documents are classified in one edge distance from the node where they were originally populated, whereas 89% of them were

---

<sup>1</sup> Precision for nodes with no documents was counted as 0.

**Table 1.** Dataset statistics and evaluation results


Dataset	Nodes	Docs	Max. subtree depth	Mi-Pr	Mi-Re	Mi-F1	Ma-Pr	Ma-Re	Ma-F1
Photography <sup>a</sup>	27	871	4	0.2218	0.1871	0.2029	0.2046	0.1165	0.1485
Beverages <sup>b</sup>	38	1456	5	0.4037	0.4938	0.4442	0.3848	0.3551	0.3693
Mammals <sup>c</sup>	88	574	5	0.3145	0.3014	0.3078	0.3854	0.2677	0.3159
Languages <sup>d</sup>	157	1217	6	0.4117	0.4503	0.4301	0.4366	0.4187	0.4275

<sup>a</sup> <http://dmoz.org/Shopping/Photography/>


<sup>b</sup> <http://dmoz.org/Shopping/Food/Beverages/>

<sup>c</sup> <http://dmoz.org/Health/Animal/Mammals/>

<sup>d</sup> [http://dmoz.org/Science/Social\\_Sciences/Linguistics/Languages/Natural/Indo-European/](http://dmoz.org/Science/Social_Sciences/Linguistics/Languages/Natural/Indo-European/)



(a)



(b)

**Fig. 5.** Analysis of the “Languages” dataset results

classified one node higher on the path to the root. Note that this still allows it to find a document of interest by browsing the classification hierarchy.

## 5 Related Work

The idea of that the get-specific classification algorithm can be encoded in a formal language and the first formal specification of the algorithm were reported in [10]. The current paper extends [10] in several respects. First, it proposes a classification model and shows how the algorithm can be implemented in this model. Second, it discusses how the model can be described and implemented in  $L^C$ . Third, it identifies the main problems peculiar to the get-specific algorithm and shows how they can be dealt with in  $L^C$ . Finally, for the first time, the current paper presents experimental results, which demonstrate that document classification can be fully automated using a knowledge-centric approach.

The idea of that natural language labels in classifications can be translated in a formal language was first introduced in [7], and, in [15], the authors provided a detailed account of the translation process using Description Logic as the target formal language. The current paper uses the translation rules described in [10], which originates from [15], but which uses the less computationally expensive propositional subset of Description Logics. In [10], the authors define a set-theoretic semantics for the translation rules and show that a propositional concept language is enough to capture the semantics of a large amount of labels.

In Information Science, hierarchical document classification usually refers to supervised or unsupervised text categorization [19]. Differently from the supervised case (e.g., see [14,9,21]), in our approach we do not need to have a pre-classified set of documents. In fact, classification choices depend on the *meaning* of classification labels and not on the documents already classified in nodes. Differently from the unsupervised approach (e.g., [16,23]), we do not need to annotate classification nodes with a relatively large (w.r.t. the label size) set of keywords to classify documents. Apart from this, in formal classification labels, the terms are connected through logical connectives, which increases the expressiveness of the category description. However, unsupervised classification is the approach closest to ours from the text categorization domain. The results, reached by the two approaches, are comparable. For instance, in [23], the authors report to reach max 42.70% in micro-F1 measure on different web directory datasets.

Noteworthy, some text categorization approaches rely on an underlying knowledge base (e.g., WordNet [17]) in order to find relations among words to optimize the construction of the feature space (e.g., see [6,18]). However, these approaches still require a training dataset to operate, i.e., they are supervised in nature.

## 6 Conclusions and Future Work

The current paper makes a contribution at the turn of several disciplines. First, it takes the notion of classification from Library Science and shows how it can be converted in a form of ontology – the fundamental notion on the Semantic Web. Interestingly, the two notions are often used interchangeably in the two communities [20]. Second, we provide a classification model and show how the get-specific algorithm, commonly used in hierarchical document classification systems, can be described in this model. Third, it shows how document classification can be fully automated using a knowledge-centric approach, an approach which is conceptually different from the one used in Information Science. Finally, evaluation results reported in this paper demonstrate the proof of concept of our approach, which makes it a viable alternative to the conventional way of automated document classification.

Our future work includes: (a) development of more accurate document concept extraction algorithms; (b) evaluation of our approach in specific domains using domain ontology as the underlying knowledge base; (c) development of knowledge base enrichment algorithms which take into account the classification semantics (which, for example, will define concept **Stater** as more specific than concept **Greek**); and (d) automatic document re-classification when the structure of the classification hierarchy changes.

## References

1. DMoz guidelines: See, <http://dmoz.org/guidelines/site-specific.html>
2. DMoz: See, <http://dmoz.org/>
3. UNSPSC: See, <http://www.unspsc.org/>

4. Yahoo! guidelines: See, <http://docs.yahoo.com/info/suggest/appropriate.html>
5. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, Cambridge (2003)
6. Bloehdorn, S., Hotho, A.: Text classification by boosting weak learners based on terms and concepts. In: *ICDM 2004*, pp. 331–334. IEEE Computer Society Press, Los Alamitos (2004)
7. Bouquet, P., Serafini, L., Zanobini, S.: Semantic coordination: a new approach and an application. In: *Proc. of the 2<sup>nd</sup> International Semantic Web Conference (ISWO'03)*. Sanibel Islands, Florida, USA (October 2003)
8. Chan, L.M., Mitchell, J.S.: *Dewey Decimal Classification: A Practical Guide*. Forest P., U.S (December 1996)
9. Dumais, S.T., Chen, H.: Hierarchical classification of web content. In: *Proc. of SIGIR-00, 23rd ACM International Conference on Research and Development in Information Retrieval*, pp. 256–263. ACM Press, Athens, GR (2000)
10. Giunchiglia, F., Marchese, M., Zaihrayeu, I.: Encoding classifications into lightweight ontologies. *JoDS VIII* (Winter 2006)
11. Giunchiglia, F., Shvaiko, P., Yatskevich, M.: Semantic schema matching. In: *Proc. of CoopIS*, pp. 347–365 (2005)
12. Giunchiglia, F., Shvaiko, P., Yatskevich, M.: Discovering missing background knowledge in ontology matching. In: *Proc. of ECAI* (2006)
13. Giunchiglia, F., Yatskevich, M.: Element level semantic matching. In: McIlraith, S.A., Plexousakis, D., van Harmelen, F. (eds.) *ISWC 2004*. LNCS, vol. 3298, Springer, Heidelberg (2004)
14. Koller, D., Sahami, M.: Hierarchically classifying documents using very few words. In: *Proc. of ICML-97, 14th International Conference on Machine Learning*, pp. 170–178. Morgan Kaufmann Publishers, Nashville (1997)
15. Magnini, B., Serafini, L., Speranza, M.: Making explicit the semantics hidden in schema models. In: Fensel, D., Sycara, K.P., Mylopoulos, J. (eds.) *ISWC 2003*. LNCS, vol. 2870, Springer, Heidelberg (2003)
16. McCallum, A., Nigam, K.: Text classification by bootstrapping with keywords, em and shrinkage. In: *Proc. of ACL99 - Workshop for Unsupervised Learning in Natural Language Processing* (1999)
17. Miller, G.: *WordNet: An electronic Lexical Database*. MIT Press, Cambridge (1998)
18. Peng, X., Choi, B.: Document classifications based on word semantic hierarchies. In: *Proc. of International Conference on Artificial Intelligence and Applications*, pp. 362–367 (2005)
19. Sebastiani, F.: Machine learning in automated text categorization. *ACM Computing Surveys* 34(1), 1–47 (2002)
20. Soergel, D.: The rise of ontologies or the reinvention of classification. *Journal of the American Society for Information Science* 50(12), 1119–1120 (1999)
21. Sun, A., Lim, E.P.: Hierarchical text classification and evaluation. In: *Proc. of ICDM*, pp. 521–528 (2001)
22. Turney, P.D.: Learning algorithms for keyphrase extraction. *Information Retrieval* 2(4), 303–336 (2000)
23. Veeramachaneni, S., Sona, D., Avesani, P.: Hierarchical dirichlet model for document classification. In: *Proc. of ICML* (2005)