

Effective use of WordNet semantics via kernel-based learning

Roberto Basili and Marco Cammisa and Alessandro Moschitti

Department of Computer Science
University of Rome "Tor Vergata", Rome, Italy
{basili,cammisa,moschitti}@info.uniroma2.it

Abstract

Research on document similarity has shown that complex representations are not more accurate than the simple *bag-of-words*. Term clustering, e.g. using latent semantic indexing, word co-occurrences or synonym relations using a word ontology have been shown not very effective. In particular, when to extend the similarity function external prior knowledge is used, e.g. WordNet, the retrieval system decreases its performance. The critical issues here are methods and conditions to integrate such knowledge.

In this paper we propose kernel functions to add prior knowledge to learning algorithms for document classification. Such kernels use a term similarity measure based on the WordNet hierarchy. The kernel trick is used to implement such space in a balanced and statistically coherent way. Cross-validation results show the benefit of the approach for the Support Vector Machines when few training data is available.

1 Introduction

The large literature on term clustering, term similarity and weighting schemes shows that document similarity is a central topic in Information Retrieval (IR). The research efforts have mostly been directed in enriching the document representation by using clustering (*term generalization*) or adding compounds (*term specifications*). These studies are based on the assumption that the similarity between two documents can be expressed as the similarity between pairs of matching terms. Following this idea,

term clustering methods based on corpus term distributions or on external prior knowledge (e.g. provided by WordNet) were used to improve the basic term matching.

An example of statistical clustering is given in (Bekkerman et al., 2001). A feature selection technique, which clusters similar features/words, called the Information Bottleneck (IB), was applied to Text Categorization (TC). Such cluster based representation outperformed the simple *bag-of-words* on only one out of the three experimented collections. The effective use of external prior knowledge is even more difficult since no attempt has ever been successful to improve document retrieval or text classification accuracy, (e.g. see (Smeaton, 1999; Sussna, 1993; Voorhees, 1993; Voorhees, 1994; Moschitti and Basili, 2004)).

The main problem of term cluster based representations seems the unclear nature of the relationship between the word and the cluster information levels. Even if (semantic) clusters tend to improve the system recall, simple terms are, on a large scale, more accurate (e.g. (Moschitti and Basili, 2004)). To overcome this problem, hybrid spaces containing terms and clusters were experimented (e.g. (Scott and Matwin, 1999)) but the results, again, showed that the mixed statistical distributions of clusters and terms impact either marginally or even negatively on the overall accuracy.

In (Voorhees, 1993; Smeaton, 1999), clusters of synonymous terms as defined in WordNet (WN) (Fellbaum, 1998) were used for document retrieval. The results showed that the misleading information due to the wrong choice of the local term senses causes the overall accuracy to decrease. Word sense disambiguation (WSD) was thus applied beforehand by indexing the documents by means of disambiguated senses, i.e. synset codes (Smeaton, 1999;

Sussna, 1993; Voorhees, 1993; Voorhees, 1994; Moschitti and Basili, 2004). However, even the state-of-the-art methods for WSD did not improve the accuracy because of the inherent noise introduced by the disambiguation mistakes. The above studies suggest that term clusters decrease the precision of the system as they force weakly related or unrelated (in case of disambiguation errors) terms to give a contribution in the similarity function. The successful introduction of prior external knowledge relies on the solution of the above problem.

In this paper, a model to introduce the semantic lexical knowledge contained in the WN hierarchy in a supervised text classification task has been proposed. Intuitively, the main idea is that the documents d are represented through the set of all pairs in the vocabulary $\langle t, t' \rangle \in V \times V$ originating by the terms $t \in d$ and all the words $t' \in V$, e.g. the WN nouns. When the similarity between two documents is evaluated, their matching pairs are used to account for the final score. The weight given to each term pair is proportional to the similarity that the two terms have in WN. Thus, the term t of the first document contributes to the document similarity according to its relatedness with any of the terms of the second document and the prior external knowledge, provided by WN, quantifies the single term to term relatedness. Such approach has two advantages: (a) we obtain a well defined space which supports the similarity between terms of different surface forms based on external knowledge and (b) we avoid to explicitly define term or sense clusters which inevitably introduce noise.

The class of spaces which embeds the above pair information may be composed by $O(|V|^2)$ dimensions. If we consider only the WN nouns (about 10^5), our space contains about 10^{10} dimensions which is not manageable by most of the learning algorithms. Kernel methods, can solve this problem as they allow us to use an implicit space representation in the learning algorithms. Among them Support Vector Machines (SVMs) (Vapnik, 1995) are kernel based learners which achieve high accuracy in presence of many irrelevant features. This is another important property as selection of the informative pairs is left to the SVM learning.

Moreover, as we believe that the prior knowledge in TC is not so useful when there is a sufficient

amount of training documents, we experimented our model in poor training conditions (e.g. less equal than 20 documents for each category). The improvements in the accuracy, observed on the classification of the well known Reuters and 20 NewsGroups corpora, show that our document similarity model is very promising for general IR tasks: unlike previous attempts, it makes sense of the adoption of semantic external resources (i.e. WN) in IR.

Section 2 introduces the WordNet-based term similarity. Section 3 defines the new document similarity measure, the kernel function and its use within SVMs. Section 4 presents the comparative results between the traditional linear and the WN-based kernels within SVMs. In Section 5 comparative discussion against the related IR literature is carried out. Finally Section 6 derives the conclusions.

2 Term similarity based on general knowledge

In IR, any similarity metric in the vector space models is driven by lexical matching. When small training material is available, few words can be effectively used and the resulting document similarity metrics may be inaccurate. Semantic generalizations overcome data sparseness problems as contributions from different but semantically similar words are made available.

Methods for the induction of semantically inspired word clusters have been widely used in language modeling and lexical acquisition tasks (e.g. (Clark and Weir, 2002)). The resource employed in most works is WordNet (Fellbaum, 1998) which contains three subhierarchies: for nouns, verbs and adjectives. Each hierarchy represents lexicalized concepts (or senses) organized according to an "is-a-kind-of" relation. A concept s is described by a set of words $syn(s)$ called *synset*. The words $w \in syn(s)$ are synonyms according to the sense s .

For example, the words *line*, *argumentation*, *logical argument* and *line of reasoning* describe a synset which expresses the methodical process of logical reasoning (e.g. "I can't follow your line of reasoning"). Each word/term may be lexically related to more than one synset depending on its senses. The word *line* is also a member of the synset *line*, *dividing line*, *demarcation* and *contrast*, as a *line* denotes

also a conceptual separation (e.g. "there is a narrow line between sanity and insanity"). The Wordnet noun hierarchy is a direct acyclic graph¹ in which the edges establish the *direct_isa* relations between two synsets.

2.1 The Conceptual Density

The automatic use of WordNet for NLP and IR tasks has proved to be very complex. First, how the topological distance among senses is related to their corresponding conceptual distance is unclear. The pervasive lexical ambiguity is also problematic as it impacts on the measure of conceptual distances between word pairs. Second, the approximation of a set of concepts by means of their generalization in the hierarchy implies a conceptual loss that affects the target IR (or NLP) tasks. For example, *black* and *white* are *colors* but are also *chess pieces* and this impacts on the similarity score that should be used in IR applications. Methods to solve the above problems attempt to map a priori the terms to specific generalizations levels, i.e. to *cuts* in the hierarchy (e.g. (Li and Abe, 1998; Resnik, 1997)), and use corpus statistics for weighting the resulting mappings. For several tasks (e.g. in TC) this is unsatisfactory: different contexts of the same corpus (e.g. documents) may require different generalizations of the same word as they independently impact on the document similarity.

On the contrary, the *Conceptual Density (CD)* (Agirre and Rigau, 1996) is a flexible semantic similarity which depends on the generalizations of word senses not referring to any fixed level of the hierarchy. The *CD* defines a metrics according to the topological structure of WordNet and can be seemingly applied to two or more words. The measure formalized hereafter adapt to word pairs a more general definition given in (Basili et al., 2004).

We denote by \bar{s} the set of nodes of the hierarchy rooted in the synset s , i.e. $\{c \in S | c \text{ isa } s\}$, where S is the set of WN synsets. By definition $\forall s \in S, s \in \bar{s}$. *CD* makes a guess about the proximity of the senses, s_1 and s_2 , of two words u_1 and u_2 , according to the information expressed by the minimal sub-hierarchy, \bar{s} , that includes them. Let S_i be the set of

generalizations for at least one sense s_i of the word u_i , i.e. $S_i = \{s \in S | s_i \in \bar{s}, u_i \in \text{syn}(s_i)\}$. The *CD* of u_1 and u_2 is:

$$CD(u_1, u_2) = \begin{cases} 0 & \text{iff } S_1 \cap S_2 = \emptyset \\ \max_{s \in S_1 \cap S_2} \frac{\sum_{i=0}^h (\mu(\bar{s}))^i}{|\bar{s}|} & \text{otherwise} \end{cases} \quad (1)$$

where:

- $S_1 \cap S_2$ is the set of WN shared generalizations (i.e. the common hypernyms) of u_1 and u_2
- $\mu(\bar{s})$ is the average number of children per node (i.e. the branching factor) in the sub-hierarchy \bar{s} . $\mu(\bar{s})$ depends on WordNet and in some cases its value can approach 1.
- h is the depth of the *ideal*, i.e. maximally dense, *tree* with enough leaves to cover the two senses, s_1 and s_2 , according to an average branching factor of $\mu(\bar{s})$. This value is actually estimated by:

$$h = \begin{cases} \lfloor \log_{\mu(\bar{s})} 2 \rfloor & \text{iff } \mu(\bar{s}) \neq 1 \\ 2 & \text{otherwise} \end{cases} \quad (2)$$

When $\mu(s)=1$, h ensures a tree with at least 2 nodes to cover s_1 and s_2 (*height* = 2).

- $|\bar{s}|$ is the number of nodes in the sub-hierarchy \bar{s} . This value is statically measured on WN and it is a negative bias for the higher level generalizations (i.e. larger \bar{s}).

CD models the semantic distance as the density of the generalizations $s \in S_1 \cap S_2$. Such *density* is the ratio between the number of nodes of the *ideal tree* and $|\bar{s}|$. The ideal tree should (a) link the two senses/nodes s_1 and s_2 with the minimal number of edges (isa-relations) and (b) maintain the same branching factor (*bf*) observed in \bar{s} . In other words, this tree provides the minimal number of nodes (and isa-relations) sufficient to connect s_1 and s_2 according to the topological structure of \bar{s} . For example, if \bar{s} has a *bf* of 2 the ideal tree connects the two senses with a single node (their father). If the *bf* is 1.5, to replicate it, the ideal tree must contain 4 nodes, i.e. the grandfather which has a *bf* of 1 and the father which has *bf* of 2 for an average of 1.5. When *bf* is 1 the Eq. 1 degenerates to the inverse of the number of nodes in the path between s_1 and s_2 , i.e. the simple proximity measure used in (Siolas and d'Alch Buc, 2000).

¹As only the 1% of its nodes own more than one parent in the graph, most of the techniques assume the hierarchy to be a tree, and treat the few exception heuristically.

It is worth noting that for each pair $CD(u_1, u_2)$ determines the similarity according to *the closest lexical senses*, $s_1, s_2 \in \bar{s}$: the remaining senses of u_1 and u_2 are irrelevant, with a resulting semantic disambiguation side effect. CD has been successfully applied to semantic tagging ((Basili et al., 2004)). As the WN hierarchies for other POS classes (i.e. verb and adjectives) have topological properties different from the noun hyponymy network, their semantics is not suitably captured by Eq. 1. In this paper, Eq. 1 has thus been only applied to noun pairs. As the high number of such pairs increases the computational complexity of the target learning algorithm, efficient approaches are needed. The next section describes how kernel methods can make practical the use of the Conceptual Density in Text Categorization.

3 A WordNet Kernel for document similarity

Term similarities are used to design document similarities which are the core functions of most TC algorithms. The term similarity proposed in Eq. 1 is valid for all term pairs of a target vocabulary and has two main advantages: (1) the relatedness of each term occurring in the first document can be computed against *all* terms in the second document, i.e. all different pairs of similar (not just identical) tokens can contribute and (2) if we use all term pair contributions in the document similarity we obtain a measure consistent with the term probability distributions, i.e. the sum of all term contributions does not penalize or emphasize arbitrarily any subset of terms. The next section presents more formally the above idea.

3.1 A semantic vector space

Given two documents d_1 and $d_2 \in D$ (the document-set) we define their similarity as:

$$K(d_1, d_2) = \sum_{w_1 \in d_1, w_2 \in d_2} (\lambda_1 \lambda_2) \times \sigma(w_1, w_2) \quad (3)$$

where λ_1 and λ_2 are the weights of the words (features) w_1 and w_2 in the documents d_1 and d_2 , respectively and σ is a term similarity function, e.g. the conceptual density defined in Section 2. To prove that Eq. 3 is a valid kernel is enough to show that it is a specialization of the general definition of convolution kernels formalized in (Haus-

ler, 1999). Hereafter, we report such definition. Let X, X_1, \dots, X_m be separable metric spaces, $x \in X$ a structure and $\vec{x} = x_1, \dots, x_m$ its parts, where $x_i \in X_i \quad \forall i = 1, \dots, m$. Let R be a relation on the set $X \times X_1 \times \dots \times X_m$ such that $R(\vec{x}, x)$ is "true" if \vec{x} are the parts of x . We indicate with $R^{-1}(x)$ the set $\{\vec{x} : R(\vec{x}, x)\}$. Given two objects x and $y \in X$ their similarity $K(x, y)$ is defined as:

$$K(x, y) = \sum_{\vec{x} \in R^{-1}(x)} \sum_{\vec{y} \in R^{-1}(y)} \prod_{i=1}^m K_i(x_i, y_i) \quad (4)$$

If X defines the document set (i.e. $D = X$), and X_1 the vocabulary of the target document corpus ($X_1 = V$), it follows that: $x = d$ (a document), $\vec{x} = x_1 = w \in V$ (a word which is a part of the document d) and $R^{-1}(d)$ defines the set of words in the document d . As $\prod_{i=1}^m K_i(x_i, y_i) = K_1(x_1, y_1)$, then $K_1(x_1, y_1) = K(w_1, w_2) = (\lambda_1 \lambda_2) \times \sigma(w_1, w_2)$, i.e. Eq. 3.

The above equation can be used in support vector machines as illustrated by the next section.

3.2 Support Vector Machines and Kernel methods

Given the vector space in \mathbb{R}^η and a set of positive and negative points, SVMs classify vectors according to a separating hyperplane, $H(\vec{x}) = \vec{\omega} \cdot \vec{x} + b = 0$, where \vec{x} and $\vec{\omega} \in \mathbb{R}^\eta$ and $b \in \mathbb{R}$ are learned by applying the *Structural Risk Minimization principle* (Vapnik, 1995). From the kernel theory we have that:

$$\begin{aligned} H(\vec{x}) &= \left(\sum_{h=1..l} \alpha_h \vec{x}_h \right) \cdot \vec{x} + b = \sum_{h=1..l} \alpha_h \vec{x}_h \cdot \vec{x} + b = \\ &= \sum_{h=1..l} \alpha_h \phi(d_h) \cdot \phi(d) + b = \sum_{h=1..l} \alpha_h K(d_h, d) + b \end{aligned} \quad (5)$$

where, d is a classifying document and d_h are all the l training instances, projected in \vec{x} and \vec{x}_h respectively. The product $K(d, d_h) = \langle \phi(d) \cdot \phi(d_h) \rangle$ is the *Semantic WN-based Kernel (SK)* function associated with the mapping ϕ .

Eq. 5 shows that to evaluate the separating hyperplane in \mathbb{R}^η we do not need to evaluate the entire vector \vec{x}_h or \vec{x} . Actually, we do not know even the mapping ϕ and the number of dimensions, η . As it is sufficient to compute $K(d, d_h)$, we can carry out the learning with Eq. 3 in the \mathbb{R}^n , avoiding to

use the explicit representation in the \mathbb{R}^n space. The real advantage is that we can consider only the word pairs associated with non-zero weight, i.e. we can use a sparse vector computation. Additionally, to have a uniform score across different document size, the kernel function can be normalized as follows:

$$\frac{SK(d_1, d_2)}{\sqrt{SK(d_1, d_1) \cdot SK(d_2, d_2)}}$$

4 Experiments

The use of WordNet (WN) in the term similarity function introduces a prior knowledge whose impact on the Semantic Kernel (SK) should be experimentally assessed. The main goal is to compare the traditional Vector Space Model kernel against SK , both within the Support Vector learning algorithm.

The high complexity of the SK limits the size of the experiments that we can carry out in a feasible time. Moreover, we are not interested to large collections of training documents as in these training conditions the simple *bag-of-words* models are in general very effective, i.e. they seem to model well the document similarity needed by the learning algorithms. Thus, we carried out the experiments on small subsets of the 20NewsGroups² (20NG) and the *Reuters-21578*³ corpora to simulate critical learning conditions.

4.1 Experimental set-up

For the experiments, we used the SVM-light software (Joachims, 1999) (available at `svmlight.joachims.org`) with the default linear kernel on the token space (adopted as the baseline evaluations). For the SK evaluation we implemented the Eq. 3 with $\sigma(\cdot, \cdot) = CD(\cdot, \cdot)$ (Eq. 1) inside SVM-light. As Eq. 1 is only defined for nouns, a part of speech (POS) tagger has been previously applied. However, also verbs, adjectives and numerical features were included in the pair space. For these tokens a $CD = 0$ is assigned to pairs made by different strings. As the POS-tagger could introduce errors, in a second experiment, any token with a successful look-up in the WN noun hierarchy was considered in the kernel. This approximation has the benefit to retrieve useful information even

²Available at `www.ai.mit.edu/people/jrennie/20Newsgroups/`.

³The Apté split available at `kdd.ics.uci.edu/databases/reuters21578/reuters21578.html`.

for verbs and capture the similarity between verbs and some nouns, e.g. *to drive* (via the noun *drive*) has a common synset with *parkway*.

For the evaluations, we applied a careful SVM parameterization: a preliminary investigation suggested that the trade-off (between the training-set error and margin, i.e. c option in SVM-light) parameter optimizes the F_1 measure for values in the range $[0.02, 0.32]$ ⁴. We noted also that the cost-factor parameter (i.e. j option) is not critical, i.e. a value of 10 always optimizes the accuracy. The feature selection techniques and the weighting schemes were not applied in our experiments as they cannot be accurately estimated from the small available training data.

The classification performance was evaluated by means of the F_1 measure⁵ for the single category and the MicroAverage for the final classifier pool (Yang, 1999). Given the high computational complexity of SK we selected 8 categories from the 20NG⁶ and 8 from the Reuters corpus⁷

To derive statistically significant results with few training documents, for each corpus, we randomly selected 10 different samples from the 8 categories. We trained the classifiers on one sample, parameterized on a second sample and derived the measures on the other 8. By rotating the training sample we obtained 80 different measures for each model. The size of the samples ranges from 24 to 160 documents depending on the target experiment.

4.2 Cross validation results

The SK (Eq. 3) was compared with the linear kernel which obtained the best F_1 measure in (Joachims, 1999). Table 1 reports the first comparative results for 8 categories of 20NG on 40 training documents. The results are expressed as the *Mean* and the *Std. Dev.* over 80 runs. The F_1 are reported in Column 2 for the linear kernel, i.e. *bow*, in Column 3 for SK without applying POS information and in Column 4

⁴We used all the values from 0.02 to 0.32 with step 0.02.

⁵ F_1 assigns equal importance to Precision P and Recall R , i.e. $F_1 = \frac{2P \cdot R}{P + R}$.

⁶We selected the 8 most different categories (in terms of their content) i.e. *Atheism*, *Computer Graphics*, *Misc Forsale*, *Autos*, *Sport Baseball*, *Medicine*, *Talk Religions* and *Talk Politics*.

⁷We selected the 8 largest categories, i.e. *Acquisition*, *Earn*, *Crude*, *Grain*, *Interest*, *Money-fx*, *Trade* and *Wheat*.

for SK with the use of POS information (SK -POS). The last row shows the MicroAverage performance for the above three models on all 8 categories. We note that SK improves bow of 3%, i.e. 34.3% vs. 31.5% and that the POS information reduces the improvement of SK , i.e. 33.5% vs. 34.3%.

To verify the hypothesis that WN information is useful in low training data conditions we repeated the evaluation over the 8 categories of Reuters with samples of 24 and 160 documents, respectively. The results reported in Table 2 shows that (1) again SK improves bow (41.7% - 37.2% = 4.5%) and (2) as the number of documents increases the improvement decreases (77.9% - 75.9% = 2%). It is worth noting that the standard deviations tend to assume high values. In general, the use of 10 disjoint training/testing samples produces a higher variability than the n -fold cross validation which insists on the same document set. However, this does not affect the t -student confidence test over the differences between the MicroAverage of SK and bow since the former has a higher accuracy at 99% confidence level.

The above findings confirm that SK outperforms the *bag-of-words* kernel in critical learning conditions as the semantic contribution of the SK recovers useful information. To complete this study we carried out experiments with samples of different size, i.e. 3, 5, 10, 15 and 20 documents for each category. Figures 1 and 2 show the learning curves for 20NG and Reuters corpora. Each point refers to the average on 80 samples.

As expected the improvement provided by SK decreases when more training data is available. However, the improvements are not negligible yet. The SK model (without POS information) preserves about 2-3% of improvement with 160 training documents. The matching allowed between noun-verb pairs still captures semantic information which is useful for topic detection. In particular, during the similarity estimation, each word activates 60.05 pairs on average. This is particularly useful to increase the amount of information available to the SVMs.

Finally, we carried out some experiments with 160 Reuters documents by discarding the string matching from SK . Only words having different surface forms were allowed to give contributions to the Eq. 3.

Category	bow	SK	SK -POS
<i>Atheism</i>	29.5±19.8	32.0±16.3	25.2±17.2
<i>Comp.Graph</i>	39.2±20.7	39.3±20.8	29.3±21.8
<i>Misc.Forsale</i>	61.3±17.7	51.3±18.7	49.5±20.4
<i>Autos</i>	26.2±22.7	26.0±20.6	33.5±26.8
<i>Sport.Baseb.</i>	32.7±20.1	36.9±22.5	41.8±19.2
<i>Sci.Med</i>	26.1±17.2	18.5±17.4	16.6±17.2
<i>Talk.Relig.</i>	23.5±11.6	28.4±19.0	27.6±17.0
<i>Talk.Polit.</i>	28.3±17.5	30.7±15.5	30.3±14.3
MicroAvg. F_1	31.5±4.8	34.3±5.8	33.5±6.4

Table 1: Performance of the linear and Semantic Kernel with 40 training documents over 8 categories of 20NewsGroups collection.

Category	24 docs		160 docs	
	bow	SK	bow	SK
<i>Acq.</i>	55.3±18.1	50.8±18.1	86.7±4.6	84.2±4.3
<i>Crude</i>	3.4±5.6	3.5±5.7	64.0±20.6	62.0±16.7
<i>Earn</i>	64.0±10.0	64.7±10.3	91.3±5.5	90.4±5.1
<i>Grain</i>	45.0±33.4	44.4±29.6	69.9±16.3	73.7±14.8
<i>Interest</i>	23.9±29.9	24.9±28.6	67.2±12.9	59.8±12.6
<i>Money-fx</i>	36.1±34.3	39.2±29.5	69.1±11.9	67.4±13.3
<i>Trade</i>	9.8±21.2	10.3±17.9	57.1±23.8	60.1±15.4
<i>Wheat</i>	8.6±19.7	13.3±26.3	23.9±24.8	31.2±23.0
Mic.Avg.	37.2±5.9	41.7±6.0	75.9±11.0	77.9±5.7

Table 2: Performance of the linear and Semantic Kernel with 40 and 160 training documents over 8 categories of the Reuters corpus.




Figure 1: MicroAverage F_1 of SVMs using bow , SK and SK -POS kernels over the 8 categories of 20NewsGroups.

The important outcome is that SK converges to a MicroAverage F_1 measure of 56.4% (compare with Table 2). This shows that the word similarity provided by WN is still consistent and, although in the worst case, slightly effective for TC: the evidence is that a suitable balancing between lexical ambiguity and topical relatedness is captured by the SVM learning.




Figure 2: MicroAverage F_1 of SVMs using *bow* and *SK* over the 8 categories of the Reuters corpus.

5 Related Work

The IR studies in this area focus on the term similarity models to embed statistical and external knowledge in document similarity.

In (Kontostathis and Pottenger, 2002) a *Latent Semantic Indexing* analysis was used for term clustering. Such approach assumes that values x_{ij} in the transformed term-term matrix represents the similarity (> 0) and anti-similarity between terms i and j . By extension, a negative value represents an anti-similarity between i and j enabling both positive and negative clusters of terms. Evaluation of query expansion techniques showed that positive clusters can improve Recall of about 18% for the *CISI* collection, 2.9% for *MED* and 3.4% for *CRAN*. Furthermore, the negative clusters, when used to prune the result set, improve the precision.

The use of external semantic knowledge seems to be more problematic in IR. In (Smeaton, 1999), the impact of semantic ambiguity on IR is studied. A WN-based semantic similarity function between noun pairs is used to improve indexing and document-query matching. However, the WSD algorithm had a performance ranging between 60-70%, and this made the overall semantic similarity not effective.

Other studies using semantic information for improving IR were carried out in (Sussna, 1993) and (Voorhees, 1993; Voorhees, 1994). Word semantic information was here used for text indexing and query expansion, respectively. In (Voorhees, 1994) it is shown that semantic information derived directly from WN without a priori WSD produces poor results.

The latter methods are even more problematic in TC (Moschitti and Basili, 2004). Word senses tend to systematically correlate with the positive examples of a category. Different categories are better characterized by different words rather than different senses. Patterns of lexical co-occurrences in the training data seem to suffice for automatic disambiguation. (Scott and Matwin, 1999) use WN senses to replace simple words without word sense disambiguation and small improvements are derived only for a small corpus. The scale and assessment provided in (Moschitti and Basili, 2004) (3 corpora using cross-validation techniques) showed that even the accurate disambiguation of WN senses (about 80% accuracy on nouns) did not improve TC.

In (Siolas and d'Alch Buc, 2000) was proposed an approach similar to the one presented in this article. A term proximity function is used to design a kernel able to semantically smooth the similarity between two document terms. Such semantic kernel was designed as a combination of the Radial Basis Function (RBF) kernel with the term proximity matrix. Entries in this matrix are inversely proportional to the length of the WN hierarchy path linking the two terms. The performance, measured over the 20NewsGroups corpus, showed an improvement of 2% over the *bag-of-words*. Three main differences exist with respect to our approach. First, the term proximity does not fully capture the WN topological information. Equidistant terms receive the same similarity irrespectively from their generalization level. For example, *Sky* and *Location* (direct hyponyms of *Entity*) receive a similarity score equal to *knife* and *gun* (hyponyms of *weapon*). More accurate measures have been widely discussed in literature, e.g. (Resnik, 1997). Second, the kernel-based *CD* similarity is an elegant combination of lexicalized and semantic information. In (Siolas and d'Alch Buc, 2000) the combination of weighting schemes, the RBF kernel and the proximity matrix has a much less clear interpretation. Finally, (Siolas and d'Alch Buc, 2000) selected only 200 features via Mutual Information statistics. In this way rare or non statistically significant terms are neglected while being source of often relevant contributions in the *SK* space modeled over WN.

Other important work on semantic kernel for retrieval has been developed in (Cristianini et al.,

2002; Kandola et al., 2002). Two methods for inferring semantic similarity from a corpus were proposed. In the first a system of equations were derived from the dual relation between word-similarity based on document-similarity and viceversa. The equilibrium point was used to derive the semantic similarity measure. The second method models semantic relations by means of a diffusion process on a graph defined by lexicon and co-occurrence information. The major difference with our approach is the use of a different source of prior knowledge. Similar techniques were also applied in (Hofmann, 2000) to derive a Fisher kernel based on a latent class decomposition of the term-document matrix.

6 Conclusions

The introduction of semantic prior knowledge in IR has always been an interesting subject as the examined literature suggests. In this paper, we used the conceptual density function on the WordNet (WN) hierarchy to define a document similarity metric. Accordingly, we defined a semantic kernel to train Support Vector Machine classifiers. Cross-validation experiments over 8 categories of 20NewsGroups and Reuters over multiple samples have shown that in poor training data conditions, the WN prior knowledge can be effectively used to improve (up to 4.5 absolute percent points, i.e. 10%) the TC accuracy.

These promising results enable a number of future researches: (1) larger scale experiments with different measures and semantic similarity models (e.g. (Resnik, 1997)); (2) improvement of the overall efficiency by exploring feature selection methods over the *SK*, and (3) the extension of the semantic similarity by a general (i.e. non binary) application of the conceptual density model.

References

E. Agirre and G. Rigau. 1996. Word sense disambiguation using conceptual density. In *Proceedings of COLING'96, Copenhagen, Denmark*.

R. Basili, M. Cammisa, and F. M. Zanzotto. 2004. A similarity measure for unsupervised semantic disambiguation. In *In Proceedings of Language Resources and Evaluation Conference*, Lisbon, Portugal.

Ron Bekkerman, Ran El-Yaniv, Naftali Tishby, and Yoad Winter. 2001. On feature distributional clustering for text categorization. In *Proceedings of SIGIR'01*, New Orleans, Louisiana, US.

Stephen Clark and David Weir. 2002. Class-based probability estimation using a semantic hierarchy. *Comput. Linguist.*, 28(2):187–206.

Nello Cristianini, John Shawe-Taylor, and Huma Lodhi. 2002. Latent semantic kernels. *J. Intell. Inf. Syst.*, 18(2-3):127–152.

Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.

D. Haussler. 1999. Convolution kernels on discrete structures. Technical report ucs-crl-99-10, University of California Santa Cruz.

Thomas Hofmann. 2000. Learning probabilistic models of the web. In *Research and Development in Information Retrieval*.

T. Joachims. 1999. Making large-scale SVM learning practical. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*.

J. Kandola, J. Shawe-Taylor, and N. Cristianini. 2002. Learning semantic similarity. In *NIPS'02) - MIT Press*.

A. Kontostathis and W. Pottenger. 2002. Improving retrieval performance with positive and negative equivalence classes of terms.

Hang Li and Naoki Abe. 1998. Generalizing case frames using a thesaurus and the mdl principle. *Computational Linguistics*, 23(3).

Alessandro Moschitti and Roberto Basili. 2004. Complex linguistic features for text classification: a comprehensive study. In *Proceedings of ECIR'04*, Sunderland, UK.

P. Resnik. 1997. Selectional preference and sense disambiguation. In *Proceedings of ACL Siglex Workshop on Tagging Text with Lexical Semantics, Why, What and How?*, Washington, 1997.

Sam Scott and Stan Matwin. 1999. Feature engineering for text classification. In *Proceedings of ICML'99*, Bled, SL. Morgan Kaufmann Publishers, San Francisco, US.

Georges Siolas and Florence d'Alch Buc. 2000. Support vector machines based on a semantic kernel for text categorization. In *Proceedings of IJCNN'00*. IEEE Computer Society.

Alan F. Smeaton. 1999. Using NLP or NLP resources for information retrieval tasks. In *Natural language information retrieval*, Kluwer Academic Publishers, Dordrecht, NL.

M. Sussna. 1993. Word sense disambiguation for free-text indexing using a massive semantic network. In *CKIM'93*.

V. Vapnik. 1995. *The Nature of Statistical Learning Theory*. Springer.

Ellen M. Voorhees. 1993. Using wordnet to disambiguate word senses for text retrieval. In *Proceedings SIGIR'93* Pittsburgh, PA, USA.

Ellen M. Voorhees. 1994. Query expansion using lexical-semantic relations. In *Proceedings of SIGIR'94*, ACM/Springer.

Y. Yang. 1999. An evaluation of statistical approaches to text categorization. *Information Retrieval Journal*.