

# Large Scale Unstructured Document Classification Using Unlabeled Data and Syntactic Information

Seong-Bae Park and Byoung-Tak Zhang

School of Computer Science and Engineering  
Seoul National University  
Seoul 151-742, Korea  
{sbpark, btzhang}@bi.snu.ac.kr

**Abstract.** Most document classification systems consider only the distribution of content words of the documents, ignoring the syntactic information underlying the documents though it is also an important factor. In this paper, we present an approach for classifying large scale unstructured documents by incorporating both lexical and syntactic information of documents. For this purpose, we use the co-training algorithm, a partially supervised learning algorithm, in which two separated views for the training data are employed and the small number of labeled data are augmented by a large number of unlabeled data. Since both lexical and syntactic information can play roles of separated views for the unstructured documents, the co-training algorithm enhances the performance of document classification using both of them and a large number of unlabeled documents. The experimental results on Reuters-21578 corpus and TREC-7 filtering documents show the effectiveness of unlabeled documents and the use of both lexical and syntactic information.

## 1 Introduction

Automatic document classification is an important research area in information retrieval and has a great potential for many applications handling text such as routing and filtering. Its aim is to assign a given document to the predefined category to which it belongs. Up to now, various kinds of algorithms based on machine learning or statistics have been applied to this problem and showed relatively high performance. However, most of them have been using a simple *bag of words* representation of documents [7]. That is, they use only the distribution of content words, assuming that the words are independent one another. But, this representation ignores *linguistic information* underlying the documents, which is another important factor in filtering texts.

Each document has its own traits in the style, and the syntactic information is one of the best measures to capture the stylistic divergence among the different kinds of documents [1]. Although the syntactic features can give much information in classifying the documents, they are not widely used due to their lack

of formal definition and complicated representation. In addition, unfortunately, the current NLP (natural language processing) techniques are not able to provide accurate results in syntax analyzing. However, some studies show that text chunks can give enough information on syntax analysis instead of full parsing [14].

Another problem in document classification is that there are a great number of inexpensive unlabeled documents while there are a few labeled documents. It is very expensive to obtain the labeled documents, since labeling of documents must be done by human experts. The co-training algorithm [2] is one of the successful algorithms handling unlabeled data. It is in general applied to the problems where there are two distinct views of each example in the dataset. For example, web pages can be considered to have two views: one is for contents of them and the other is for link information. It learns separate classifiers over each of the views, and augments a small number of labeled examples incorporating unlabeled examples.

In order to incorporate unlabeled examples with labeled examples, a metric to measure the confidence of labeling an unlabeled example is required in the co-training algorithm. And, recently the concept of margin is introduced with the emergence of Support Vector Machine [13], where the margin is defined as the product of true label (assumed to be -1 or +1) and the classifier's output. If the prediction of the classifier is correct, the margin is positive, otherwise negative. When the output of the classifier is real value, it can be considered as a confidence on prediction.

In this paper, we propose a co-trained Support Vector Machines (SVM) for document classification. It is based on the co-training algorithm, so that it effectively uses not only given small number of labeled documents but a great number of unlabeled documents. For the two views of the co-training algorithm, we use both lexical information and syntactic information. Thus, the proposed method can be applied to classify the unstructured normal documents.

This paper is organized as follows. Section 2 explains several approaches to incorporating unlabeled documents and linguistic information in document classification. Section 3 describes the co-training algorithm using both the lexical and the syntactic information. Section 4 reports the experimental results on Reuters-21578 corpus and TREC-7 filtering documents. Finally, Section 5 draws conclusions.

## 2 Related Work

Incorporating unlabeled examples is not a new scheme in information retrieval. It has been a major research topic under the name of *pseudo relevance feedback*. The idea of relevance feedback is to examine a portion of the retrieved documents and assign relevance values to each of them. The contents of the documents are analyzed to make the query closer towards the examined relevant documents and away from the examined non-relevant documents. The most common model for relevance feedback is proposed by Rocchio [11]. The terms in the new query

are ranked by the weighted sum of the term weights in the current query, the known relevant documents, and the known non-relevant documents. That is, the new query is constructed with several top terms ranked by:

$$w_i(Q') = \alpha \cdot w_i(Q) + \beta \cdot w_i(R) + \gamma \cdot w_i(N),$$

where  $w_i(Q')$  and  $w_i(Q)$  are the weights of a term  $w_i$  in the new and current queries,  $w_i(R)$  and  $w_i(N)$  are the weights of  $w_i$  in the relevant and non-relevant documents, and  $\alpha$ ,  $\beta$ , and  $\gamma$  are the parameters to be determined. Many experiments on relevance feedback empirically show that it improves retrieval performance dramatically [12]. Pseudo relevance feedback typically adds new terms to the initial query by assuming that the several top documents in the initial ranked output are relevant.

In most discussions of relevance feedback, it is usually assumed that the initial query is of high quality, so that a major portion of the returned documents by the initial query are highly relevant. However, in many situations such an assumption is impractical. In document classification, we always can not trust all the labels of unlabeled documents estimated by the current classifier, since usually there are a small number of labeled documents while there are a great number of unlabeled documents. If the trained classifier does not coincide to the intrinsic model which generates the documents, the performance will be hurt by the unlabeled documents. Thus, more sophisticated methods are required for document classification using unlabeled ones.

One way to overcome the limit of traditional term weighting method is to use syntactic and semantic information. The merits of syntactic information is investigated by Hull [6]. Other studies also show that it is quite useful as content-identifier [15]. They state that using non-NLP generated phrases as terms in vector space retrieval is helpful at low recall level, while it is not helpful at high recall level. However, it is required to obtain more accurate syntactic information than just predefined window-size word sequences for its general use.

Because the current NLP techniques do not provide accurate information enough to be used in information retrieval, text chunking is considered to be an alternative to full parsing [14]. Text chunking is to divide text into syntactically related non-overlapping segments of words. Since this task is formulated as estimating an identifying function from the information (features) available in the surrounding context, various techniques have been applied to it [4].

### 3 Co-training Algorithm for Classifying Unstructured Documents

#### 3.1 Co-training Algorithm

The co-training algorithm is one of the successful algorithms handling unlabeled examples [2]. It is in general applied to the problems where there are two distinct views of each example in the dataset. It learns separate classifiers over each of the views, and augments a small set of labeled examples by incorporating

Given  $L$ : a set of labeled examples, and  $U$ : a set of unlabeled examples.

**Do Until** there is no unlabeled example.

1. Train classifier  $h_1$  on view  $V_1$  of  $L$ .
2. Train classifier  $h_2$  on view  $V_2$  of  $L$ .
3. Allow  $h_1$  to determine labels of examples in  $U$ .
4. Allow  $h_2$  to determine labels of examples in  $U$ .
5. Determine  $U' \subset U$ , whose elements are most confidently labeled by  $h_1$  and  $h_2$ .
6.  $U = U \setminus U'$
7.  $L = L + U'$

**Fig. 1.** An abstract of the co-training algorithm.

unlabeled examples. Its final prediction is made by combining their predictions to decrease classification errors. The larger is the variance of the classifiers when both classifiers are unbiased, the better is the performance of the algorithm. Since the co-training uses two classifiers with distinct views, its performance is better than any single classifier.

Figure 1 outlines the co-training algorithm. It uses two distinct views  $V_1$  and  $V_2$  when learning from labeled and unlabeled data, and incrementally upgrades classifiers ( $h_1$  and  $h_2$ ) over each view. Each classifier is initialized with a few labeled examples. At every iteration, each classifier chooses unlabeled examples and adds them to the set of labeled examples,  $L$ . The selected unlabeled examples are those which each classifier can determine their labels with the highest confidence. After that, the classifiers are trained again using the augmented labeled set. The final output of the algorithm is given as a combination of the two classifiers. Given an example  $\mathbf{x}$  to be classified, the probability of the possible class  $c_j$  is determined by multiplying two posterior probabilities. The class  $c^*$  of  $\mathbf{x}$  is set to the one with the highest probability:

$$c^* = \arg \max_{c_j \in C} (P(c_j | \mathbf{x}) = P_{h_1}(c_j | \mathbf{x})P_{h_2}(c_j | \mathbf{x})),$$

where  $C$  is the set of all possible classes.

Blum and Mitchell formalized the co-training settings and provided theoretical guarantee under certain assumptions [2]. Their first assumption is that the data distribution is compatible with the target function. That is, the target functions over each view predict the same class for most examples. The second assumption is that the views are conditionally independent with each other. If this assumption holds, the added examples at each iteration will be at least as informative as random examples. Thus, the iteration can progress though there are some mislabeled examples during the learning. However, these assumptions are somewhat unrealistic in practice, since the views from the same data are inclined to be related with each other in some way. Nigam and Ghani performed thorough empirical investigation on the conditional independence of the views [10]. Their experiments showed that the view independence affects the performance of the co-training algorithm, but the algorithm is still more effective than

other algorithms incorporating unlabeled examples even when some dependence exist between the views.

For the classifiers in the co-training algorithm, we adopt Support Vector Machines that show significant improvement over other machine learning algorithms when applied to document classification [7]. At each iteration of the co-training, the most confident  $|U'|$  examples should be selected from  $U$ . The SVMs provide a natural way to calculate the confidence of labeling unlabeled examples by a margin. The margin  $m$  for an unlabeled example  $\mathbf{x}_i$  is defined as

$$m = y_i(\mathbf{w} \cdot \mathbf{x}_i + b),$$

where  $y_i \in \{-1, +1\}$  is the label predicted by the hyperplane with the trained parameters  $\mathbf{w}$  and  $b$  of SVMs. That implies, the margin can be considered to be a distance from  $\mathbf{x}_i$  to the hyperplane, assuming that the predicted label is correct. The more distant  $\mathbf{x}_i$  lies from the hyperplane, the more confident it is to predict the label of  $\mathbf{x}_i$ . Since SVMs are not probabilistic models, the final prediction of the co-training can not be made by multiplying the probabilities of each classifier. Instead, the final prediction is made only by the classifier whose margin is larger than the other.

### 3.2 Two Views

One possible view for document classification is to treat each document as a vector whose elements are the weight to the vocabulary. Most machine learning algorithms applied to document classification adopt this representation. The main drawbacks of this representation are that (i) it assumes that each word in the documents is independent one another, and (ii) it ignores much linguistic information underlying in the documents.

Stamatatos et al. showed experimentally that the syntactic information among various kinds of linguistic information is a reliable clue for document classification [14]. One additional benefit in using syntactic information for document classification by the co-training algorithm is that it is somewhat independent from term weights. Unfortunately, the current natural language processing techniques are not able to provide accurate syntactic analysis results. However, the *text chunks* instead of full parsing are good features enough to provide syntactic information for document classification. The chunks can be obtained in high accuracy with superficial investigation.

Therefore, we can define two distinct views for unstructured documents, so that the co-training algorithm can be naturally applied to classifying them. The two views are:

#### – Lexical Information

Most machine learning algorithm applied to automatic document classification are based on  $tf \cdot idf$ , a commonly used term weighting scheme in information retrieval. The  $tf$  factor is an estimation of the occurrence probability of a term if it is normalized, and the  $idf$  is the amount of information related to the occurrence of the term.

### – Syntactic Information

Each document is represented in a vector in which the elements are syntactic features, and the features are derived from text chunking. This information can support finding particular or specific style of the documents.

### 3.3 Syntactic Features

To represent the documents in vectors whose elements are syntactic information, all documents need to be represented with chunk information. Since the unstructured documents are normally raw, all sentences in the documents must be chunked in the preprocessing step of classification. The purpose of chunking is to divide a given sentence into nonoverlapping segments. Let us consider the following sentence.

*Example 1.* He reckons the current deficit will narrow to only # 1.8 billion in September .

This sentence, composed of 15 words including a period, is grouped into nine segments after chunking as follows:

*Example 2.* [NP He] [VP reckons] [NP the current deficit] [VP will narrow] [PP to] [NP only # 1.8 billion ] [PP in] [NP September] [O .]

In order to chunk the sentences in the documents, the lexical information and the POS (part-of-speech) information on the contextual words are required. Brill’s tagger [3] is used to obtain POS tags for each word in the documents. The chunk type of each word is determined by Support Vector Machines trained with the dataset of CoNLL-2000 shared task<sup>1</sup>.

Although there are 12 types of phrases in CoNLL-2000 dataset, we consider, in this paper, only five types of phrases: NP<sup>2</sup>, VP, ADVP, PP, and O, where O implies none of NP, VP, ADVP, and PP. This is reasonable because those five types take 97.31% of the dataset and are major phrases for constituting a sentence. Each phrase except O has two kinds of chunk types: B-XP and I-XP. For instance, B-NP represents the first word of the noun phrase, while I-NP is given to other words in noun phrase. Thus, we consider 11 chunk types.

The use of Support Vector Machines showed the best performance in the shared task [8]. The contexts used to identify the chunk type of the  $i$ -th word  $w_i$  in a sentence are:

$$w_j, POS_j \quad (j = i - 2, i - 1, i, i + 1, i + 2)$$

$$c_j \quad (j = 1 - 2, i - 1)$$

where  $POS_j$  and  $c_j$  are respectively the POS tag and the chunk type of  $w_i$ . Since SVMs are basically binary classifiers and there are 11 types of chunks, SVMs are extended to multi-class classifiers by *pairwise classification* [13].

<sup>1</sup> <http://lcs-www.uia.ac.be/conll2000/chunking>

<sup>2</sup> NP represents a noun phrase, VP a verb phrase, ADVP an adverb phrase, and PP a prepositional phrase.

**Table 1.** Syntactic features for document classification.

Feature	Description
SF1	detected NPs / total detected chunks
SF2	detected VPs / total detected chunks
SF3	detected PPs / total detected chunks
SF4	detected ADVPs / total detected chunks
SF5	detected Os / total detected chunks
SF6	words included in NPs / detected NPs
SF7	words included in VPs / detected VPs
SF8	words included in PPs / detected PPs
SF9	words included in ADVPs / detected ADVPs
SF10	words included in Os / detected Os
SF11	sentences / words

Table 1 shows the features used to represent documents using text chunks. Top five features represent how often the phrases are used in a document, the following five features imply how long they are, and the final feature means how long a sentence is on the average. That is, every document is represented in a 11-dimensional vector.

## 4 Experiments

### 4.1 Datasets

**Reuters-21578.** The Reuters-21578 corpus is the most commonly used benchmark corpus in text classification. It consists of over 20,000 the Reuters newswire articles in the period between 1987 to 1991, and has 135 kinds of topics while only 10 of them are used for experiments. There are three versions to divide this corpus into a training set and a test set: “ModLewis”, “ModApte”, and “Mod-Hayes”. Among them, “ModApte” which is most widely used is employed in this paper. In this version, there are 9603 training documents, 3299 test documents, and 27863 unique words after stemming and stop word removal.

**TREC-7.** The dataset for the TREC-7 filtering track consists of the articles from AP News for the years between 1988 to 1990, Topics 1-50. The 1988 AP documents are used for a training set, and the 1989-1990 AP documents for a test set. We do not use any information from the topic, such as ‘description’. For documents representation using lexical information, we stem them and remove words from the stop-list. No thesaurus or other datasets are used. Though there are more than 160,000 words in this dataset, we choose only 7289 most frequently appearing words. There are 79,898 articles in the training set, but only 9,572 articles, just about 12% of them are labeled. Thus, we regard those 9,572 articles as a labeled document set, and the remaining 88% of the articles as an unlabeled document set.

**Table 2.** The accuracy improvement by using syntactic information on Reuters-21578 corpus. *Lexical* implies when we use only  $tf \cdot idf$ , *Syntactic* when we use only text chunks, and *Both* when we use both of them. The performance measure is LF1.

Class	Relevant Documents	Lexical	Syntactic	Both
earn	2877	2876	2504	<b>2895</b>
acq	1650	1587	1028	<b>1642</b>
money-fx	538	188	147	<b>193</b>
grain	433	<b>26</b>	14	<b>26</b>
crude	389	292	150	<b>312</b>
trade	369	153	114	<b>155</b>
interest	347	130	112	<b>141</b>
wheat	212	113	88	<b>116</b>
ship	197	98	70	<b>108</b>
corn	181	86	68	<b>87</b>

## 4.2 Performance Measure

To evaluate the classification performance of the proposed method, we use *utility* measure, for example LF1. Let  $R_+$  be the number of relevant and retrieved documents,  $N_+$  the number of non-relevant but retrieved documents,  $R_-$  the number of relevant but non-retrieved documents, and  $N_-$  the number of non-relevant and non-retrieved documents. Then, *linear utility* is defined as

$$\text{Linear Utility} = aR_+ + bN_+ + cR_- + dN_-,$$

where  $a, b, c$  and  $d$  are constant coefficients. The F1 and F3 measures are defined:

$$\text{LF1} = 3R_+ - 2N_+,$$

$$\text{LF2} = 3R_+ - N_+.$$

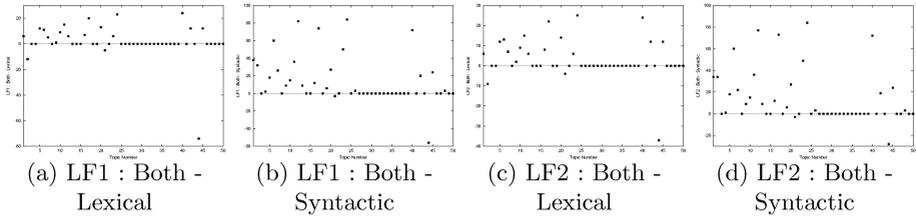
The drawback of linear utility is that it is meaningless to average through all topics because a few topics will have dominant effects on the results. Thus, to show the whole performance effectively we use the *scaled utility* [5] defined as

$$\text{Scaled Utility} = \frac{\max\{u(S, T), U(s)\} - U(s)}{\text{Max}U(T) - U(s)},$$

where  $u(S, T)$  is the utility of system  $S$  for topic  $T$ ,  $\text{Max}U(T)$  is the maximum possible utility score for topic  $T$ , and  $U(s)$  is the utility of retrieving  $s$  non-relevant documents. Since we consider only 10 topics in Reuters-21578 corpus, we use LF1 measure. However, there are 50 topics in TREC-7, so that we use the scaled linear utility.

## 4.3 Effect of Syntactic Information

**Reuters-21578.** When we do not consider unlabeled examples on Reuters-21578 dataset, the effect of using additional syntactic information is given in



**Fig. 2.** The performance improvement by using both lexical and syntactic information. Two different measures are used: LF1 and LF2.

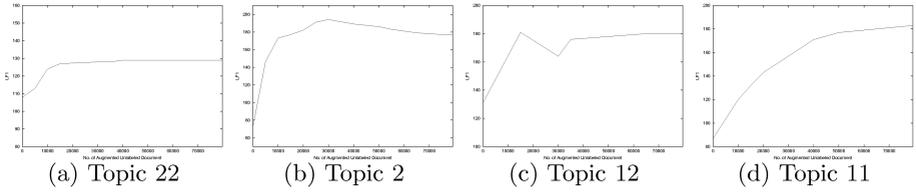
**Table 3.** The result on TREC-7 dataset. The performance measure is the averaged scaled LF1 and LF2 when  $U(s)$  is set to 0.

Measure	Lexical	Syntactic	Both
LF1	0.2005	0.1680	<b>0.2192</b>
LF2	0.2010	0.2010	<b>0.2155</b>

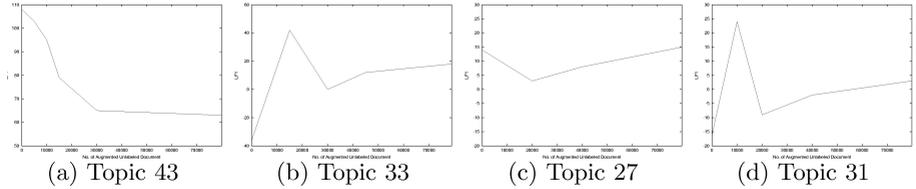
Table 2. The classification performance using both lexical and syntactic information outperforms the one using any single information for most classes. We find that the larger is the number of relevant documents in the training set, the higher is the LF1 for using lexical information. This is because the SVMs tend to achieve lower error rate and higher utility value when a class has a large number of relevant documents. This trend is kept except *grain* class even when we use both information. For *grain* class, the extremely low LF1 for the syntactic information causes the performance using both information to be dominated by the lexical information.

**TREC-7.** Figure 2 shows the advantage of using both lexical and syntactic information. These results are obtained in considering both labeled and unlabeled documents. The X-axis is the topic numbers, while the Y-axis is the difference of performance that we obtain by using both kinds of information rather than single information. Thus, positive values in these graphs mean that using both information outperforms using single one. We find by these figures that using both kinds of information improves the performance for both LF1 and LF2. In addition, the difference between *Both-Syntactic* is larger than *Both-Lexical* for LF1. This result coincides with Table 2 that using only lexical information is better than using only syntactic information.

The averaged performance of Figure 2 is given in Table 3. Since the results in the table did not enter TREC pools, all the unjudged documents in the test set are ignored. The use of both information gives the best result for both measures. Using only lexical information is better than using only syntactic information in LF1, but achieves the same performance in LF2. This coincides with the work of Mitra et al. [9] that the performance of a phrase-based retrieval system is not superior than that of a term-based retrieval system. Even though we do not use the phrases for the index of documents, a similar result is obtained that



**Fig. 3.** The analysis of the effects of unlabeled documents for the topics with a great number of relevant documents in TREC-7. The performance measure is LF1.

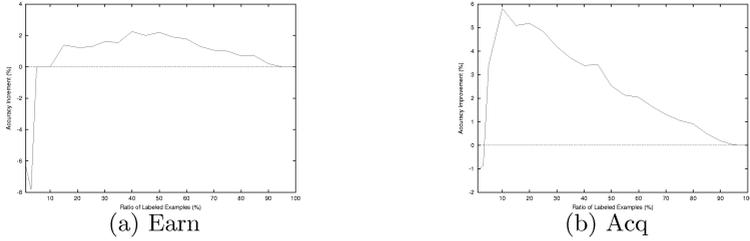


**Fig. 4.** The analysis of the effects of unlabeled documents for the topics with a small number of relevant documents in TREC-7. The performance measure is LF1.

the information derived from phrases are not better than terms. We have not studied this phenomenon in greater details, but we presume three reasons for it. First, we do not consider any content words in the document for syntactic information. Thus, it overgeneralizes the documents with styles. Second, there are only 1,468 reporters in the training set of TREC-7, though there are more than 79,000 articles. One reporter should write about 50 articles on the average, but more than a half of them write less than 10 articles. That is, a few reporters wrote a large portion of the articles. This forces the syntactic character of each reporter not to reveal and the performance of using only syntactic information to be low. The other reason is that there are a great number of non-sentence forms such as tables and listings in the training set since they are news articles. They make the syntactic analysis failed.

#### 4.4 Effect of Unlabeled Documents

**TREC-7.** Figure 3 and Figure 4 show that the performance of the proposed method is enhanced by unlabeled documents even though there are a small number of relevant documents. The figures in Figure 3 are on the topics with a great number of relevant documents, while the figures in Figure 4 are on the topics with a small number of relevant documents. For example, *Topic 22* has 358 relevant documents, but *Topic 31* has only 1 relevant ones. Most topics except *Topic 43* show the similar trends that the performance gets higher by using unlabeled documents and gets higher and higher as more unlabeled documents are used in the learning. This is a good point of Support Vector Machines since they tend to correctly classify the relevancy with the larger number of documents.



**Fig. 5.** The improvement in accuracy by using additional unlabeled documents.

**Reuters-21578.** Many previous work asserted that unlabeled documents improve the performance of document classification [16]. Our experiments also show that unlabeled documents are useful for better performance. Figure 5 shows the effectiveness of unlabeled documents involved in the co-training algorithm. The X-axis represents the ratio of labeled documents to total documents, while Y-axis is the accuracy improvement by additional unlabeled documents. For *earn*, the unlabeled documents play a positive role when more than 10% of training documents are labeled. So is for *acq*, when more than 7% of training documents are labeled.

However, even when we obtain the highest improvement by unlabeled documents, it does not reach to the best performance when we know the labels of all the training examples beforehand. For example, the improvement of accuracy is 5.81% when 10% of documents are labeled in *acq*. In this case, the accuracy is just 89.93% while the accuracy with 100% labeled documents is 95.21%. This implies that some of the unlabeled documents are mislabeled during the co-training process and have a negative effect on the classifiers. The effectiveness of unlabeled documents can be maximized when the number of labeled ones is small. To fill a gap of this difference, human intervention is needed. But, it is still an open problem when to intervene in the process.

## 5 Conclusion

We proposed the co-trained Support Vector Machines for document filtering. This method uses both the traditional lexical information and the syntactic information for the co-training algorithm and makes the algorithm applied to unstructured document classification, not only web page classification. In addition, with the algorithm we can incorporate ubiquitous unlabeled documents naturally to augment a small number of given labeled documents. The experiments on Reuters-21578 corpus and TREC-7 filtering documents show that using the syntactic information with the traditional lexical information improves the classification performance and the unlabeled documents are good resources to overcome the limited number of labeled documents. This is important because we can construct the classifier for large scale documents in high performance

with just a small number of labeled documents. While the effectiveness of unlabeled documents is empirically proved, another problem is revealed that we need a method to overcome the mislabeling during the co-training algorithm.

**Acknowledgement.** This research was supported by the Korean Ministry of Education under the BK21-IT program and by the Korean Ministry of Science and Technology under BrainTech program.

## References

1. D. Biber. *Dimensions of Register Variation: A Cross-Linguistic Comparison*. Cambridge University Press, 1995.
2. A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of CONLT-98*, pages 92–100, 1998.
3. E. Brill. A simple rule-based part-of-speech tagger. In *Proceedings of ANLP-92*, pages 152–155, 1992.
4. CoNLL. *Shared Task for Computational Natural Language Learning (CoNLL)*. <http://lcg-www.uia.ac.be/conll2000/chunking>, 2000.
5. D. Hull. The TREC-7 filtering track: Description and analysis. In *Proceedings of TREC-7*, pages 33–56, 1998.
6. D. Hull, G. Grefenstette, B. Schulze, E. Gaussier, H. Schutze, and J. Pedersen. Xerox TREC-5 site report: Routing, filtering, nlp, and spanish tracks. In *Proceedings of TREC-7*, 1997.
7. T. Joachims. Text categorization with support vector machines: Learning with many relevant features. In *Proceedings of ECML-98*, pages 137–142, 1998.
8. T. Kudo and Y. Matsumoto. Use of support vector learning for chunk identification. In *Proceedings of CoNLL-2000 and LLL-2000*, pages 142–144, 2000.
9. M. Mitra, C. Buckley, A. Singhal, and C. Cardie. An analysis of statistical and syntactic phrases. In *Proceedings of RIAO-97*, pages 200–214, 1997.
10. K. Nigam and R. Ghani. Analyzing the effectiveness and applicability of Co-training. In *Proceedings of CIKM-2000*, pages 86–93, 2000.
11. J. Rocchio. Relevance feedback in information retrieval. *The SMART Retrieval System: Experiments in Automatic Document Processing*, pages 313–323, 1971.
12. G. Salton and C. Buckley. Improving retrieval performance by relevance feedback. *Journal of the American Society for Information Science*, pages 288–297, 1990.
13. B. Scholkopf, C. Burges, and A. Smola. *Advances in Kernel Methods: Support Vector Machines*. MIT Press, 1999.
14. E. Stamatatos, N. Fakotakis, and G. Kokkinakis. Automatic text categorization in terms of genre and author. *Computational Linguistics*, 26(4):471–496, 2000.
15. A. Turpin and A. Moffat. Statistical phrases for vector-space information retrieval. In *Proceedings of SIGIR-1999*, pages 309–310, 1999.
16. T. Zhang and F. Oles. A probability analysis on the value of unlabeled data for classification problems. In *Proceedings of ICML-2000*, pages 1191–1198, 2000.