Late-Breaking Developments in the Field of Artificial Intelligence
Papers Presented at the Twenty-Seventh AAAI Conference on Artificial Intelligence

# Cotraining Based Bilingual Sentiment Lexicon Learning

**Dehong Gao[†*], Furu Wei[‡], Wenjie Li[†], Xiaohua Liu[‡] and Ming Zhou[‡]**
[†]The Hong Kong Polytechnic University, Hong Kong, China
[‡]Microsoft Research Asia, Beijing, China
[†]{csdgao,cswjli}@comp.polyu.edu.hk; [‡]{fuwei,xiaoliu,mingzhou}@microsoft.com

## Abstract

In this paper, we address the issue of bilingual sentiment lexicon learning(**BSLL**) which aims to automatically and simultaneously generate sentiment words for two languages. The underlying motivation is that sentiment information from two languages can perform iterative mutual-teaching in the learning procedure. We propose to develop two classifiers to determine the sentiment polarities of words under a co-training framework, which makes full use of the two-view sentiment information from the two languages. The word alignment derived from the parallel corpus is leveraged to design effective features and to bridge the learning of the two classifiers. The experimental results on English and Chinese languages show the effectiveness of our approach in BSLL.

## Introduction

Sentiment lexicon is regarded as the most valuable resource for sentiment analysis (Pang and Lee 2008). Extensive research has been conducted on automatic sentiment lexicon learning to avoid manually annotating sentiment words (Qiu et al. 2011), while these research mainly focuses on English sentiment lexicon learning (Turney and Littman 2003) and may not work well in other languages due to the lack of necessary resources in these languages.

In this paper, we focus on *bilingual sentiment lexicon learning*, which aims to automatically and simultaneously generate sentiment lexicons for two languages. The underlying motivation is that the sentiment information available in two languages can be interactively used to enhance the learning process of both languages. In order to bridge the two learning processes, the words in two languages are connected based on bilingual resources. Specially, for a pair of the connected words, two different views of features are collected to indicate their sentiment polarities from the perspectives of their own languages. We formalize bilingual sentiment lexicon learning as a word polarity classification task and propose a co-training based approach to take advantage of these two views of features. We develop a classifier for each language to distinguish its words with positive, negative or neutral sentiment polarity and then generate bilingual sentiment lexicons based on the classification results. The benefit of the proposed co-training approach is that the sentiment information from large amounts of the unlabeled data can be incrementally used to improve the classification performance.

We choose English as one of the languages in BSLL because of its rich available resources. We can collect many sentiment words from public available English sentiment lexicons as training data. Meanwhile, since the co-training process can be advanced in the interactions of the sentiment information between two languages, we are allowed to only collect a small set of training data for the other language and transfer the rich sentiment information from English to the other language during the iterations. This can greatly reduce the effort on training data collection. In addition, we leverage the word alignment derived from a parallel corpus to design effective features and to build relations between words in English and the other language. We develop two classifiers to predict the polarities of all the unlabeled words in the parallel corpus to generate bilingual sentiment lexicons. The experimental results on Chinese and English sentiment lexicon learning show that the proposed co-training approach outperforms inductive and transductive approaches.

## Co-training for BSLL

In this paper we formalize BSLL as a word sentiment classification task. Given the labeled words (with sentiment polarities) and unlabeled words in two languages, the task is to develop two classifiers for the two languages. The two classifiers can be used to predict sentiment polarities (i.e. positive, negative and neutral) of all the unlabeled words. We can select the top confident ones to build the bilingual sentiment lexicons. Specially, we develop the two classifiers ($C_E$ and $C_T$) under the co-training framework which effectively leverages the sentiment information from two languages. For each language we collect one set of sentiment words as labeled dataset ($L_E$/$L_T$) and one set of words whose sentiment polarities are unknown, as unlabeled dataset ($U_E$/$U_T$). To bridge the learning processes of the two languages, we build word relations $R$ between words in two languages based on bilingual resources. Thus for a pair of connected words, we can obtain two views of features (i.e. $F_E$ and $F_T$) to indicate their sentiment polarity. The proposed co-training approach is to train two classifiers based on these two-view features. In each iteration, the classifiers predict the sentiment polarities of all words in unlabeled datasets and then the most-confident ones are added to their corresponding labeled datasets for the next iteration of training.

[*]Contribution during internship at Microsoft Research Asia

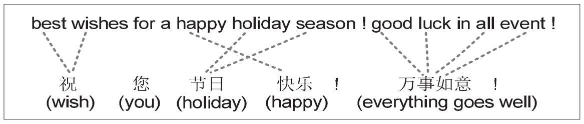


Figure 1: Parallel corpus and word alignment

Meanwhile, according to the most-confident words in one language, we select the confident words for the other language based on the bilingual word relations and then use these newly-selected words to update the classifier of the other languages as well. In the final classification, the two classifiers are used to predict the sentiment polarities of all the words in unlabeled datasets and to generate the bilingual sentiment lexicons. The detail co-training algorithm for BSLL is described in Algorithm 1.

**Algorithm 1** The co-training algorithm

**Input:** Labeled/unlabeled datasets $L_E$/$L_T$ and $U_E$/$U_T$ for English/the other language. Features $F_E$ and $F_T$ derived from bilingual resources. Parameters $p$, $q$ and $N$.
**Output:** Classifiers $C_E$ and $C_T$
1: **while** $k < N$ **do**
2: Train classifier $C_E$ from $L_E$ based on $F_E$
3: Use $C_E$ to label the polarities of words in $U_E$
4: Select top-$p$ positive/negative/neutral of the most confidently words $S_E$ from $U_E$
5: Similar to Line2∼4, train classifier $C_T$ for the other language
6: Select top-$q$ words $A_T$/$A_E$ from $S_E$/$S_T$ based on $R$
7: Remove $S_E$/$S_T$ and $A_E$/$A_T$ from $U_E$/$U_T$, and add them to $L_E$/$L_T$, respectively
8: $k = k + 1$
9: **end while**

## Word Relation Building and Feature Extraction Based on Parallel Corpus

To overcome the language gap, we use the huge amount of parallel sentences, which are used as the foundation of the state-of-the-art statistical machine translation engines. As shown in Figure 1, the two sentences in English and Chinese are parallel sentences, which express the same meaning in different languages. We can then obtain the word alignment information from the sentence pairs, like the Chinese word “” (happy) linked to the English word “*happy*” and we say these two words are aligned. The distinct advantages of using parallel corpus in BSLL are 1) the parallel corpus can be extracted from web documents and the extracted words and expressions are closer to the human language. Likewise, the word alignment extracted from the parallel sentences is more suitable for learning sentiment words which tend to be colloquial. 2) It allows us to explore a new kind of classification features, i.e. word alignment features.

Based on parallel corpus, we extract one list of aligned words in the other language for a given word. These aligned words can be regarded as the synonyms of the given word in the other language. Intuitively two words may share the same sentiment polarity if they are aligned together by parallel sentences. We assume that the top-$q$ aligned words can keep the same sentiment polarity with their original word. Based on this assumption, if one word is confidently predicted by one classifier, we use the top-$q$ aligned words to update the training dataset of the other classifier. In addition, it is also reasonable to assume that if two words in a language are approximately aligned with some same words, these two words may also share the same sentiment polarity. We can use such word alignment information as classification features. It is worth to mention that, we have actually also examined the roles of some additional linguistic features like synonyms and word definitions etc in our empirical studies. However, they don't add values to the learning performance on the top of the word alignment features.

Table 1: Evaluation of bilingual sentiment lexicons

| | Positive | | | Negative | | |
|---|---|---|---|---|---|---|
| *P@N* | 100 | 1K | 5K | 100 | 1K | 5K |
| SVM(CN) | 0.68 | 0.43 | 0.38 | 0.69 | 0.43 | 0.42 |
| TSVM(CN) | 0.72 | 0.49 | 0.46 | 0.71 | 0.49 | 0.47 |
| Co-training(CN) | **0.79** | **0.55** | **0.51** | **0.77** | **0.54** | **0.52** |
| SVM(EN) | 0.74 | 0.72 | 0.51 | 0.70 | 0.51 | 0.40 |
| TSVM(EN) | 0.76 | 0.73 | 0.55 | 0.76 | 0.60 | 0.45 |
| Co-training(EN) | **0.80** | **0.76** | **0.61** | **0.81** | **0.66** | **0.53** |

## Experimental Evaluation

We experiment on English and Chinese sentiment lexicon learning and select the sentiment words in *General Inquirer* (Stone 1997) and *OpinionFinder Lexicon* (Wilson, Wiebe, and Hoffmann 2005) as the labeled English data. Words in the ISI Chinese-English parallel corpus (Munteanu and Marcu 2005) are collected as the unlabeled dataset. *BerkeleyAligner* (Liang, Taskar, and Klein 2006) is applied to align the parallel sentences in the ISI parallel corpus. The proposed co-training approach is compared with the following four baseline approaches. ***SVM(EN/CN)***: These two approaches use the inductive SVM, *LibSVM*, for English/Chinese word polarity classification based on the English/Chinese features. ***TSVM(EN/CN)***: These two approaches use the transductive SVM, for English/Chinese word polarity classification based on the English/Chinese features. Since $SVM^{light}$ cannot perform multi-class classification directly, we employ one-class-against-the-rest classification for this task (Joachims 2002).

We rank the words which are finally predicted as positive or negative according to the classifier outputs and regard these two lists are bilingual sentiment lexicons. We evaluate the quality of the bilingual sentiment lexicons with $P@N$. For $P@5K$, we annotate the sentiment polarities of 1$K$ words randomly selected from the top-5$K$ words to reduce the labor-intensive human annotation. In Table 1, the parameters of $p$, $q$, $N$ are set to 10, 3 and 100, respectively, which are tuned empirically. We find that the co-training approach remarkably outperforms the baseline approaches.

## Conclusion and Future Work

In this paper we present our preliminary study and some brief experimental results on BSLL. In the future, we will continue applying our co-training model to explore other language pairs using the parallel corpus available in the machine translation community.

## Acknowledgments

The work described in this paper was supported by a Hong Kong RGC project (PolyU No. 5202/12E) and a National Nature Science Foundation of China (NSFC No. 61272291).

## References

Joachims, T. 2002. *Learning to Classify Text Using Support Vector Machines*. Kluwer Academic Publishers.

Liang, P.; Taskar, B.; and Klein, D. 2006. Alignment by agreement. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, 104–111.

Munteanu, D. S., and Marcu, D. 2005. Improving machine translation performance by exploiting non-parallel corpora. *Comput. Linguist.* 31(4):477–504.

Pang, B., and Lee, L. 2008. *Opinion mining and sentiment analysis*, volume 2. Foundations and Trends in Information Retrieval.

Qiu, G.; Liu, B.; Bu, J.; and Chen, C. 2011. Opinion word epansion and target extraction through double propagation. *Association for Computational Linguistic* 37(1):9–27.

Stone, P. J. 1997. Thematic text analysis: New agendas for analyzing text content. *Text Analysis for the Social Sciences*.

Turney, P. D., and Littman, M. L. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems* 21(4):315–346.

Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, HLT '05, 347–354. Stroudsburg, PA, USA: Association for Computational Linguistics.