

Cross-lingual Sentiment Lexicon Learning With Bilingual Word Graph Label Propagation

Dehong Gao^{†,*}

The Hong Kong Polytechnic University

Furu Wei^{**}

Microsoft Research

Wenjie Li[†]

The Hong Kong Polytechnic University

Xiaohua Liu^{**}

Microsoft Research

Ming Zhou^{**}

Microsoft Research

In this article we address the task of cross-lingual sentiment lexicon learning, which aims to automatically generate sentiment lexicons for the target languages with available English sentiment lexicons. We formalize the task as a learning problem on a bilingual word graph, in which the intra-language relations among the words in the same language and the inter-language relations among the words between different languages are properly represented. With the words in the English sentiment lexicon as seeds, we propose a bilingual word graph label propagation approach to induce sentiment polarities of the unlabeled words in the target language. Particularly, we show that both synonym and antonym word relations can be used to build the intra-language relation, and that the word alignment information derived from bilingual parallel sentences can be effectively leveraged to build the inter-language relation. The evaluation of Chinese sentiment lexicon learning shows that the proposed approach outperforms existing approaches in both precision and recall. Experiments conducted on the NTCIR data set further demonstrate the effectiveness of the learned sentiment lexicon in sentence-level sentiment classification.

* Contribution during internship at Microsoft Research (Beijing).

** Furu Wei, Xiaohua Liu, and Ming Zhou are from Microsoft Research, Beijing, China.

E-mail: {fuwei, xiaoliu, mingzhou}@microsoft.com.

† Dehong Gao and Wenjie Li are from the Department of Computing, the Hong Kong Polytechnic University. E-mail: gaodehong_polyu@163.com, cswjli@comp.polyu.edu.hk.

Submission received: 28 April 2013; revised submission received: 25 February 2014; accepted for publication: 12 May 2014.

doi:10.1162/COLLA_00207

1. Introduction

A sentiment lexicon is regarded as the most valuable resource for sentiment analysis (Pang and Lee 2008), and lays the groundwork of much sentiment analysis research, for example, sentiment classification (Yu and Hatzivassiloglou 2003; Kim and Hovy 2004) and opinion summarization (Hu and Liu 2004). To avoid manually annotating sentiment words, an automatically learning sentiment lexicon has attracted considerable attention in the community of sentiment analysis. The existing work determines word sentiment polarities either by the statistical information (e.g., the co-occurrence of words with predefined sentiment seed words) derived from a large corpus (Riloff, Wiebe, and Wilson 2003; Hu and Liu 2006) or by the word semantic information (e.g., synonym relations) found in existing human-created resources (e.g., WordNet) (Takamura, Inui, and Okumura 2005; Rao and Ravichandran 2009). However, current work mainly focuses on English sentiment lexicon generation or expansion, while sentiment lexicon learning for other languages has not been well studied.

In this article, we address the issue of cross-lingual sentiment lexicon learning, which aims to generate sentiment lexicons for a non-English language (hereafter referred to as “the target language”) with the help of the available English sentiment lexicons. The underlying motivation of this task is to leverage the existing English sentiment lexicons and substantial linguistic resources to label the sentiment polarities of the words in the target language. To this end, we need an approach to transferring the sentiment information from English words to the words in the target language. The few existing approaches first build word relations between English and the target language. Then, based on the word relation and English sentiment seed words, they determine the sentiment polarities of the words in the target language. In these two steps, relation-building plays a fundamental role because it is responsible for the transfer of sentiment information between the two languages. Two approaches are often used to connect the words in different languages in the literature. One is based on translation entries in cross-lingual dictionaries (Hassan et al. 2011). The other relies on a machine translation (MT) engine as a black box to translate the sentiment words in English to the target language (Steinberger et al. 2011). The two approaches in Duh, Fujino, and Nagata (2011) and Mihalcea, Banea, and Weibe (2007) tend to use a small set of vocabularies to translate the natural language, which leads to a low coverage of generated sentiment lexicons for the target language.

To solve this problem, we propose a generic approach to addressing the task of cross-lingual sentiment lexicon learning. Specifically, we model this task with a bilingual word graph, which is composed of two **intra-language** subgraphs and an **inter-language** subgraph. The intra-language subgraphs are used to model the semantic relations among the words in the same languages. When building them, we incorporate both synonym and antonym word relations in a novel manner, represented by positive and negative sign weights in the subgraphs, respectively. These two intra-language subgraphs are then connected by the inter-language subgraph. We propose Bilingual word graph Label Propagation (BLP), which simultaneously takes the inter-language relations and the intra-language relations into account in an iterative way. Moreover, we leverage the word alignment information derived from a *parallel corpus* to build the inter-language relations. We connect two words from different languages that are aligned to each other in a parallel sentence pair. Taking advantage of a large parallel corpus, this approach significantly improves the coverage of the generated sentiment lexicon. The experimental results on Chinese sentiment lexicon learning show the effectiveness of the proposed approach in terms of both precision and recall. We further

evaluate the impact of the learned sentiment lexicon on sentence-level sentiment classification. When using words in the learned sentiment lexicon as features for sentiment classification of the target language, the sentiment classification can achieve a high performance.

We make the following contributions in this article.

1. We present a generic approach to automatically learning sentiment lexicons for the target language with the available sentiment lexicon in English, and we formalize the cross-lingual sentiment learning task on a bilingual word graph.
2. We build a bilingual word graph by using synonym and antonym word relations and propose a bilingual word graph label propagation approach, which effectively leverages the inter-language relations and both types (synonym and antonym) of the intra-language relations in sentiment lexicon learning.
3. We leverage the word alignment information derived from a large number of parallel sentences in sentiment lexicon learning. We build the inter-language relation in the bilingual word graph upon word alignment, and achieve significant results.

2. Related Work

2.1 English Sentiment Lexicon Learning

In general, the work on sentiment lexicon learning focuses mainly on English and can be categorized as co-occurrence-based approaches (Hatzivassiloglou and McKeown 1997; Riloff, Wiebe, and Wilson 2003; Qiu et al. 2011) and semantic-based approaches (Mihalcea, Banea, and Wiebe 2007; Takamura, Inui, and Okumura 2005; Kim and Hovy 2004).

The **co-occurrence-based approaches** determine the sentiment polarity of a given word according to the statistical information, like the co-occurrence of the word to pre-defined sentiment seed words or the co-occurrence to product features. The statistical information is mainly derived from certain corpora. One of the earliest work conducted by Hatzivassiloglou and McKeown (1997) assumes that the conjunction words can convey the polarity relation of the two words they connect. For example, the conjunction word *and* tends to link two words with the same polarity, whereas the conjunction word *but* is likely to link two words with opposite polarities. Their approach only considers adjectives, not nouns or verbs, and it is unable to extract adjectives that are not conjoined by conjunctions. Riloff et al. (2003) define several pattern templates and extract sentiment words by two bootstrapping approaches. Turney and Littman (2003) calculate the pointwise mutual information (PMI) of a given word with positive and negative sets of sentiment words. The sentiment polarity of the word is determined by average PMI values of the positive and negative sets. To obtain PMI, they provide queries (consisting of the given word and the sentiment word) to the search engine. The number of hits and the position (if the given word is near the sentiment word) are used to estimate the association of the given word to the sentiment word. Hu and Liu (2004) research sentiment word learning on customer reviews and they assume that the sentiment words tend to be correlated with product features. The frequent nouns and noun phrases are treated as product features. Then they extract the adjective words

as sentiment words from those sentences that contain one or more product features. This approach may work on a product review corpus, where one product feature may frequently appear. But for other corpora, like news articles, this approach may not be effective. Qiu et al. (2011) combine sentiment lexicon learning and opinion target extraction. A double propagation approach is proposed to learn sentiment words and to extract opinion targets simultaneously, based on eight manually defined rules.

The **semantic-based approaches** determine the sentiment polarity of a given word according to the word semantic relation, like the synonyms of sentiment seed words. The word semantic relation is usually obtained from dictionaries, for example, *WordNet*.¹ Kim and Hovy (2004) assume that the synonyms of a positive (negative) word are positive (negative) and its antonyms are negative (positive). Initializing with a set of sentiment words, they expand sentiment lexicons based on these two kinds of word relations. Kamps et al. (2004) build a synonym graph according to the synonym relation (synset) derived from *WordNet*. The sentiment polarity of a word is calculated by the shortest path to two sentiment words *good* and *bad*. However, the shortest path cannot precisely describe the sentiment orientation, considering there are only five steps between the word *good* and the word *bad* in *WordNet* (Hassan et al. 2011). Takamura et al. (2005) construct a word graph with the gloss of *WordNet*. Words are connected if a word appears in the gloss of another. The word sentiment polarity is determined by the weight of its connections on the word graph. Based on *WordNet*, Rao and Ravichandran (2009) exploit several graph-based semi-supervised learning methods like Mincuts and Label Propagation. The word polarity orientations are induced by initializing some sentiment seed words in the *WordNet* graph. Esuli et al. (2006, 2007) and Baccianella et al. (2010) treat sentiment word learning as a machine learning problem, that is, to classify the polarity orientations of the words in *WordNet*. They select seven positive words and seven negative words and expand them through the *see-also* and *antonym* relations in *WordNet*. These expanded words are then used for training. They train a ternary classifier to predict the sentiment polarities of all the words in *WordNet* and use the glosses (textual definitions of the words in *WordNet*) as the features of classification. The sentiment lexicon generated is the well-known *SentiWordNet*.²

2.2 Cross-Lingual Sentiment Lexicon Learning

The work on cross-lingual sentiment lexicon learning is still at an early stage and can be categorized into two types, according to how they bridge the words in two languages.

Mihalcea et al. (2007) generate sentiment lexicon for Romanian by directly translating the English sentiment words into Romanian through bilingual English–Romanian dictionaries. When confronting multiword translations, they translate the multiwords word by word. Then the validated translations must occur at least three times on the Web. The approach proposed by Hassan et al. (2011) learns sentiment words based on English *WordNet* and *WordNets* in the target languages (e.g., Hindi and Arabic). Cross-lingual dictionaries are used to connect the words in two languages and the polarity of a given word is determined by the average hitting time from the word to the English sentiment word set. These approaches connect words in two languages based on cross-lingual dictionaries. The main concern of these approaches is the effect of morphological inflection (i.e., a word may be mapped to multiple words in cross-lingual dictionaries).

1 <http://wordnet.princeton.edu/>.

2 <http://sentiwordnet.isti.cnr.it/>.

For example, one single English word typically has four Spanish or Italian word forms (two each for gender and for number) and many Russian word forms (due to gender, number, and case distinctions) (Steinberger et al. 2011). Usually, this approach requires an additional process to disambiguate the sentiment polarities of all the morphological variants.

To improve the sentiment classification for the target language, Banea, Mihalcea, and Wiebe (2010) translate the English sentiment lexicon into the target language using Google Translator.³ Similarly, Google Translator is used by Steinberger et al. (2011). They manually produce two high-level gold-standard sentiment lexicons for two languages (e.g., English and Spanish) and then translate them into the third language (e.g., Italian) via Google Translator. They believe that those words in the third language that appear in both translation lists are likely to be sentiment words. These approaches connect the words in two languages based on MT engines. The main concern of these approaches is the low overlapping between the vocabularies of natural documents and the vocabularies of the documents translated by MT engines (Duh, Fujino, and Nagata 2011; Meng et al. 2012a). The shortcoming of these MT-based approaches inevitably leads to low coverage.

Our task resembles the task of cross-lingual sentiment classification, like Wan (2009), Lu et al. (2011), and Meng et al. (2012a), which classifies the sentiment polarities of product reviews. Generally, these studies use semi-supervised learning approaches and regard translations from labeled English sentiment reviews as the training data. The terms in each review are leveraged as the features for training, which has proven to be effective in sentiment classification (Pang and Lee 2008). We can regard the task of sentiment lexicon learning as word-level sentiment classification. However, for word-level sentiment classification, it is not straightforward to extract features for a single word. Without sufficient features, it is difficult for these approaches to perform well in learning. Another line of cross-lingual sentiment classification uses Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003) or its variants, like Boyd-Graber and Resnik (2010) or He, Alani, and Zhou (2010). These studies assume that each review is a mixture of sentiments and each sentiment is a probability over words. Then they apply the LDA-like approach to model the sentiment polarity of each review. Nonetheless, this assumption may not be applicable in sentiment lexicon learning because a single word can be regarded as the minimal semantic unit, and it is difficult, if not impossible, to infer the latent topics from a single word. Recall that different from the sentiment classification of product reviews where the instances are normally independent, words in sentiment lexicon learning are highly related with each other, like synonyms and antonyms. Through these relations, the words can naturally form a word graph. Thus we use the graph-based learning approach to leverage the word distributions in sentiment lexicon learning. In the next section, we will introduce our proposed graph-based cross-lingual sentiment lexicon learning.

3. Cross-Lingual Sentiment Lexicon Learning

In this work, we model the task of cross-lingual sentiment lexicon learning with a bilingual word graph, where (1) the words in the two languages are represented by the nodes in two intra-language subgraphs, respectively; (2) the synonym and antonym word relations within each language are represented by the positive and negative sign

³ <http://translate.google.com/>.

weights in the corresponding intra-language subgraphs; and (3) the two intra-language subgraphs are connected by an inter-language subgraph. Mathematically, we build a graph $\mathcal{G} = (\mathbf{X}_E \cup \mathbf{X}_T, \mathbf{W}_E \cup \tilde{\mathbf{W}}_E \cup \mathbf{W}_T \cup \tilde{\mathbf{W}}_T \cup \mathbf{W}_A)$ that consists of two intra-language subgraphs $\mathcal{G}_E = (\mathbf{X}_E, \mathbf{W}_E \cup \tilde{\mathbf{W}}_E)$ and $\mathcal{G}_T = (\mathbf{X}_T, \mathbf{W}_T \cup \tilde{\mathbf{W}}_T)$ as shown in Figure 1. These two subgraphs are connected by the inter-language graph $\mathcal{G}_R = (\mathbf{X}_E \cup \mathbf{X}_T, \mathbf{W}_A)$. The elements of \mathbf{W}_E , \mathbf{W}_T , and \mathbf{W}_A are positive real numbers, that is, \mathbf{W}_E , \mathbf{W}_T , and $\mathbf{W}_A \in \mathbb{R}_+$, and $\tilde{\mathbf{W}}_E$ and $\tilde{\mathbf{W}}_T \in \mathbb{R}_-$. Because \mathcal{G} incorporates the words in two languages, we call it a **Bilingual Word Graph**. Specifically, the positive weights, \mathbf{W}_E and \mathbf{W}_T , represent the synonym intra-language relations, and the negative weights, $\tilde{\mathbf{W}}_E$ and $\tilde{\mathbf{W}}_T$, represent the antonym intra-language relations. The inter-language relations, \mathbf{W}_A , represent the connections between the words in the two languages. For cross-lingual sentiment lexicon learning, $\mathbf{X}_E = \{\mathbf{X}_E^L, \mathbf{X}_E^U\}$ denotes the labeled and unlabeled words in English and \mathbf{X}_T denotes the unlabeled words in the target language. Given the labels $\mathbf{Y}_E^L = \{y_{E_1}, \dots, y_{E_l}\}$ of the seeds \mathbf{X}_E^L , we aim to predict the sentiment polarities of the words \mathbf{X}_T . In the remainder of this section, we will present the bilingual word graph construction and the algorithm of bilingual word graph label propagation.

3.1 Bilingual Word Graph Building With Parallel Corpus and Word Alignment

We represent the words in English and in the target language as the nodes of the bilingual word graph. We use the synonym and antonym relations of the words in the same language to build \mathbf{W} and $\tilde{\mathbf{W}}$ in the intra-language graph, respectively. In the rest of this section, we will focus on how to build the inter-language relation.

Intuitively, there are two ways to connect the words in two languages. One is to insert links to the words if there exist entry mappings between the words in bilingual dictionaries (e.g., the English–Chinese dictionary). This method is simple and straightforward, but it suffers from two limitations. (1) Dictionaries are static during a certain period, whereas the sentiment lexicon evolves over time. (2) The entries in dictionaries tend to be the expressions of formal and written languages, but people prefer using colloquial language in expressing their sentiments or opinions on-line. These limitations lead to the low coverage of the links from English to the target language. An alternative way is to use an MT engine as a black box to build the inter-language relation. One can send each word in English to a publicly available MT engine and get the translations in the target language. Edges can then be inserted into the graph between the English

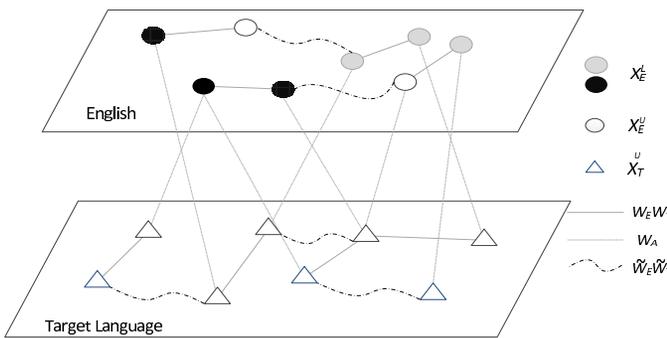


Figure 1
Bilingual word graph for cross-lingual sentiment lexicon learning.



Figure 2 Parallel sentences with word alignments.

words and their corresponding translations. This approach suffers from the problem of low coverage as well because MT engines tend to use a small set of vocabularies (Duh, Fujino, and Nagata 2011).

In this article, we propose to leverage a large bilingual parallel corpus, which is readily available in the MT research community, to build the bilingual word graph. The parallel corpus consists of a large number of parallel sentence pairs from two different languages that have been used as the foundation of the state-of-the-art statistical MT engines. Like the example shown in Figure 2, the two sentences in English and Chinese are parallel sentences, which express the same meaning in different languages. We can easily derive the word alignment from the sentence pairs, automatically using a state-of-the-art toolkit, like *GIZA++*⁴ or *BerkeleyAligner*.⁵ In this example, the Chinese word 快乐 (happy) is linked to the English word *happy* and we say that these two words are aligned. Similarly the English words *best* and *wishes* are both aligned to 祝 (wish).

The word alignment information encodes the rich association information between the words from the two languages. We are therefore motivated to leverage the parallel corpus and word alignment to build the bilingual word graph for cross-lingual sentiment lexicon learning. We take the words from both languages in the bilingual parallel corpus as the nodes in the bilingual word graph, and build the inter-language relations by connecting the two words that are aligned together in a sentence pair from a parallel corpus. There are several advantages of using a parallel corpus to build the inter-language subgraph. First, large parallel corpora are extensively used for training statistical MT engines and can be easily reused in our task. The parallel sentence pairs are usually automatically collected and mined from the Web. As a result, they contain the different and practical variations of words and phrases embedded in sentiment expressions. Second, the parallel corpus can be dynamically changed when necessary because it is relatively easy to collect from the Web. Consequentially, the novel sentiment information inferred from the parallel corpus can easily update the existing sentiment lexicons. These advantages can greatly improve the coverage of the generated sentiment lexicon, as demonstrated later in our experiments.

3.2 Bilingual Word Graph Label Propagation

As commonly used semi-supervised approaches, label propagation (Zhu and Ghahramani 2002) and its variants (Zhu, Ghahramani, and Lafferty 2003; Zhou et al. 2004) have been applied to many applications, such as part-of-speech tagging (Das and Petrov 2011; Li, Graca, and Taskar 2012), image annotation (Wang, Huang, and Ding 2011), protein function prediction (Jiang 2011; Jiang and McQuay 2012), and so forth.

4 <http://www.statmt.org/ Moses/giza/GIZA++.html>.

5 <http://nlp.cs.berkeley.edu>.

The underlying idea of label propagation is that the connected nodes in the graph tend to share the same sentiment labels. In bilingual word graph label propagation, the words tend to share same sentiment labels if they are connected by synonym relations or word alignment and tend to belong to different sentiment labels if connected by antonym relations.

In this article we propose bilingual word graph label propagation for cross-lingual sentiment lexicon learning. Let $\mathbf{F} = \{\mathbf{F}_E, \mathbf{F}_T\}$ denote the predicted labels of the unlabeled words \mathbf{X} . The **loss function** can be defined as

$$\mathbf{E}_l(\mathbf{F}) = \mu \sum_{i=1}^n \|f_{E_i} - y_{E_i}\|^2 + \mu \sum_{i=1}^m \|f_{T_i} - y_{T_i}\|^2 \quad (1)$$

where n and m denote the numbers of English words and words in the target language. Let $\mathbf{Y} = \{\mathbf{Y}_E, \mathbf{Y}_T\}$ denote the initial sentiment labels of all the words; the loss function means that the prediction could not change too much from the initial label assignment. Similar to Zhou et al. (2004), we define the **smoothness function** to indicate that if two words are connected by synonym relation or by word alignment, then they tend to share the same sentiment label. The smoothness function can be further represented by two parts, that is, the inter-language smoothness $\mathbf{E}_s^{inter}(\mathbf{F})$ and the synonym intra-language smoothness $\mathbf{E}_s^{intra}(\mathbf{F})$

$$\mathbf{E}_s^{inter}(\mathbf{F}) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m w_{A_{ij}} \left\| \frac{f_{E_i}}{\sqrt{d_{AL_{ii}}}} - \frac{f_{T_j}}{\sqrt{d_{AR_{jj}}}} \right\|^2 \quad (2)$$

$$\mathbf{E}_s^{intra}(\mathbf{F}) = \rho_1 \frac{1}{2} \sum_{i,j=1}^n w_{E_{ij}} \left\| \frac{f_{E_i}}{\sqrt{d_{E_{ii}}}} - \frac{f_{E_j}}{\sqrt{d_{E_{jj}}}} \right\|^2 + \rho_2 \frac{1}{2} \sum_{i,j=1}^m w_{T_{ij}} \left\| \frac{f_{T_i}}{\sqrt{d_{T_{ii}}}} - \frac{f_{T_j}}{\sqrt{d_{T_{jj}}}} \right\|^2 \quad (3)$$

\mathbf{D}_{AL} and \mathbf{D}_{AR} are defined as $\mathbf{D}_{AL} = \text{diag}(\sum_j w_A(1,j), \dots, \sum_j w_A(n,j))'$ and $\mathbf{D}_{AR} = \text{diag}(\sum_i w_A(i,1), \dots, \sum_i w_A(i,m))'$. \mathbf{D}_E and \mathbf{D}_T are the degree matrices of the synonym intra-language relations \mathbf{W}_E and \mathbf{W}_T , respectively. We then define the **distance function** to indicate that if two words are connected by the antonym relation they tend to belong to different sentiment labels. The distance function can be defined as

$$\mathbf{E}_d^{intra}(\mathbf{F}) = \rho_3 \frac{1}{2} \sum_{i,j=1}^n |\tilde{w}_{E_{ij}}| \left\| \frac{f_{E_i}}{\sqrt{\tilde{d}_{E_{ii}}}} - \frac{f_{E_j}}{\sqrt{\tilde{d}_{E_{jj}}}} \right\|^2 + \rho_4 \frac{1}{2} \sum_{i,j=1}^m |\tilde{w}_{T_{ij}}| \left\| \frac{f_{T_i}}{\sqrt{\tilde{d}_{T_{ii}}}} - \frac{f_{T_j}}{\sqrt{\tilde{d}_{T_{jj}}}} \right\|^2 \quad (4)$$

where $\tilde{\mathbf{D}}_E$ and $\tilde{\mathbf{D}}_T$ are the degree matrices of the absolute value of the antonym intra-language relations $\tilde{\mathbf{W}}_E$ and $\tilde{\mathbf{W}}_T$, respectively. Intuitively, for the inter-language smoothness and the synonym intra-language smoothness, the nearer the words connect with each other, the better performance will be achieved, whereas for the antonym intra-language distance, the farther the better. The objective functions can be defined as

$$\arg \min(\mathbf{E}(\mathbf{F})) = \arg \min(\mathbf{E}_s^{intra}(\mathbf{F}) + \mathbf{E}_s^{inter}(\mathbf{F}) + \mathbf{E}_l(\mathbf{F}))$$

$$\arg \max(\mathbf{E}(\mathbf{F})) = \arg \max(\mathbf{E}_d^{intra}(\mathbf{F}))$$

Thus, we define the whole objective function for cross-lingual sentiment lexicon learning as

$$\arg \min(\mathbf{E}(\mathbf{F})) = \arg \min(\mathbf{E}_s^{intra}(\mathbf{F}) + \mathbf{E}_s^{inter}(\mathbf{F}) + \mathbf{E}_l(\mathbf{F}) - \mathbf{E}_d^{intra}(\mathbf{F})) \quad (5)$$

To obtain the solution to Equation (5), we differentiate the objective function according to \mathbf{F}_E and \mathbf{F}_T , and we have

$$\frac{\partial \mathbf{E}(\mathbf{F})}{\partial \mathbf{F}_E} \Big|_{\mathbf{F}_E = \mathbf{F}_E^*} = \rho_1 \mathbf{S}_E \mathbf{F}_E + \frac{1}{2} \mathbf{S}_A \mathbf{F}_T - \rho_3 \tilde{\mathbf{S}}_E \mathbf{F}_E + \mu \mathbf{F}_E - \mu \mathbf{Y}_E = 0 \quad (6)$$

$$\frac{\partial \mathbf{E}(\mathbf{F})}{\partial \mathbf{F}_T} \Big|_{\mathbf{F}_T = \mathbf{F}_T^*} = \rho_2 \mathbf{S}_T \mathbf{F}_T + \frac{1}{2} \mathbf{S}_A' \mathbf{F}_E - \rho_4 \tilde{\mathbf{S}}_T \mathbf{F}_T + \mu \mathbf{F}_T - \mu \mathbf{Y}_T = 0 \quad (7)$$

where \mathbf{P}' is the transpose of the matrix \mathbf{P} . The graph Laplacians \mathbf{S}_E and \mathbf{S}_T of the synonym intra-language relations are $(\mathbf{I} - \mathbf{D}_E^{-\frac{1}{2}} \mathbf{W}_E \mathbf{D}_E^{-\frac{1}{2}})$ and $(\mathbf{I} - \mathbf{D}_T^{-\frac{1}{2}} \mathbf{W}_T \mathbf{D}_T^{-\frac{1}{2}})$, where \mathbf{I} is the identity matrix. The graph Laplacians $\tilde{\mathbf{S}}_E$ and $\tilde{\mathbf{S}}_T$ of the antonym intra-language relations are $(\mathbf{I} - \tilde{\mathbf{D}}_E^{-\frac{1}{2}} \tilde{\mathbf{W}}_E \tilde{\mathbf{D}}_E^{-\frac{1}{2}})$ and $(\mathbf{I} - \tilde{\mathbf{D}}_T^{-\frac{1}{2}} \tilde{\mathbf{W}}_T \tilde{\mathbf{D}}_T^{-\frac{1}{2}})$, which has been proven to be positive semi-definite (Kunegis et al. 2010). The graph Laplacian \mathbf{S}_A of the inter-language relation is $(\mathbf{I} - \mathbf{D}_{AL}^{-\frac{1}{2}} \mathbf{W}_A \mathbf{D}_{AR}^{-\frac{1}{2}})$. From Equations (6) and (7), we can obtain the optimal solutions

$$(\mathbf{M}_E - \mathbf{S}_A \mathbf{M}_T^{-1} \mathbf{S}_A') \mathbf{F}_E = 2\mu(\mathbf{Y}_E - \mathbf{S}_A \mathbf{M}_T^{-1} \mathbf{Y}_T) \quad (8)$$

$$(\mathbf{M}_T - \mathbf{S}_A' \mathbf{M}_E^{-1} \mathbf{S}_A) \mathbf{F}_T = 2\mu(\mathbf{Y}_T - \mathbf{S}_A' \mathbf{M}_E^{-1} \mathbf{Y}_E) \quad (9)$$

where $\mathbf{M}_E = 2\rho_1 \mathbf{S}_E - 2\rho_3 \tilde{\mathbf{S}}_E + 2\mu \mathbf{I}$ and $\mathbf{M}_T = 2\rho_2 \mathbf{S}_T - 2\rho_4 \tilde{\mathbf{S}}_T + 2\mu \mathbf{I}$. To avoid computing the inverse matrix in Equations (8) and (9), we apply the Jacobi algorithm (Saad 2003) to calculate the solutions as described in Algorithm 1. In line 1, we set the label of the positive seed x_i as $y_{E+i}^L = (1, 0)$ and the label of the negative seed x_j as $y_{E-j}^L = (0, 1)$. We set the label \mathbf{Y}_E^U of the unlabeled words as zero, and then generate \mathbf{Y}_E with \mathbf{Y}_E^L and \mathbf{Y}_E^U . Line 2 sets \mathbf{Y}_T as zero matrix. In line 3, we compute the matrixes $\mathbf{S}_E, \tilde{\mathbf{S}}_E, \mathbf{S}_T, \tilde{\mathbf{S}}_T, \mathbf{S}_A$, and then compute the matrixes \mathbf{M}_E and \mathbf{M}_T . The sentiment information is simultaneously

Algorithm 1. Bilingual word graph label propagation

Input: Given $\mathcal{G} = (\mathbf{X}_E \cup \mathbf{X}_T, \mathbf{W}_E \cup \tilde{\mathbf{W}}_E \cup \mathbf{W}_T \cup \tilde{\mathbf{W}}_T \cup \mathbf{W}_A), \mathbf{X}_E$,
label \mathbf{Y}_E^L for \mathbf{X}_E^L , initialize μ and $\rho_{1 \sim 4}$

Output: \mathbf{F}_T for \mathbf{X}_T and \mathbf{F}_E for \mathbf{X}_E

1. Initialize \mathbf{Y}_E with the English sentiment seeds
 2. Set \mathbf{Y}_T as zero
 3. Calculate $\mathbf{S}_E, \tilde{\mathbf{S}}_E, \mathbf{S}_T, \tilde{\mathbf{S}}_T$, and \mathbf{S}_A , then calculate \mathbf{M}_E and \mathbf{M}_T
 4. Loop
 5. $f_{E_i}^{(t+1)} = \frac{1}{(\mathbf{M}_E - \mathbf{S}_A \mathbf{M}_T^{-1} \mathbf{S}_A')_{ii}} (2\mu(\mathbf{Y}_E - \mathbf{S}_A \mathbf{M}_T^{-1} \mathbf{Y}_T)_i - \sum_{j \neq i} (\mathbf{M}_E - \mathbf{S}_A \mathbf{M}_T^{-1} \mathbf{S}_A')_{ij} f_{E_j}^{(t)})$
 6. $f_{T_i}^{(t+1)} = \frac{1}{(\mathbf{M}_T - \mathbf{S}_A' \mathbf{M}_E^{-1} \mathbf{S}_A)_{ii}} (2\mu(\mathbf{Y}_T - \mathbf{S}_A' \mathbf{M}_E^{-1} \mathbf{Y}_E)_i - \sum_{j \neq i} (\mathbf{M}_T - \mathbf{S}_A' \mathbf{M}_E^{-1} \mathbf{S}_A)_{ij} f_{T_j}^{(t)})$
 7. Until \mathbf{F}_E and \mathbf{F}_T convergence
-

propagated through lines 4–7 until the predicted labels F_E and F_T are converged. For an unlabeled word x_i , if $|f(i, 0) - f(i, 1)| < \xi$ (ξ is set as $1.0E^{-4}$), x_i is regarded as neutral; if $(f(i, 0) - f(i, 1)) \geq \xi$, x_i is regarded as positive; and if $(f(i, 1) - f(i, 0)) \geq \xi$, x_i is regarded as negative.

4. Experiment

4.1 Data Sets

We conduct experiments on Chinese sentiment lexicon learning. As in previous work (Baccianella, Esuli, and Sebastiani 2010), the sentiment words in General Inquirer lexicon are selected as the English seeds (Stone 1997). From the GI lexicon we collect 2,005 positive words and 1,635 negative words. To build the bilingual word graph, we adopt the Chinese–English parallel corpus, which is obtained from the news articles published by *Xinhua* News Agency in Chinese and English collections, using the automatic parallel sentence identification approach (Munteanu and Marcu 2005). Altogether, we collect more than 25M parallel sentence pairs in English and Chinese. We remove all the stopwords in Chinese and English (e.g., 的 (of) and *am*) together with the low-frequency words that occur fewer than 5 times. After preprocessing, we finally have more than 174,000 English words, among which 3,519 words have sentiment labels and more than 146,000 Chinese words for which we need to predict the sentiment labels. To transfer sentiment information to Chinese unlabeled words more efficiently, we remove the unlabeled English words in the word graph (i.e., $X_E^U = \Phi$). The unsupervised method, namely, *BerkeleyAligner*, is used to align the parallel sentences in this article (Liang, Taskar, and Klein 2006). As an unsupervised method, it does not require us to manually collect training data and does not need the complex training processing, and its performance is competitive with supervised methods. With these two advantages, we can focus more on our task of cross-lingual sentiment lexicon learning. Based on the word alignment derived by *BerkeleyAligner*, the inter-language W_A is initialized with the normalized alignment frequency. The English and Chinese versions⁶ of WordNet are used to build the intra-language relations W_E , \tilde{W}_E , W_T , and \tilde{W}_T , respectively. WordNet (Miller 1995) groups words into synonym sets, called synsets. We collect about 117,000 synsets from the English WordNet and about 80,000 synsets from the Chinese WordNet. In total, we obtain 8,406 and 6,312 antonym synset pairs.

We first generate both positive and negative scores for each unlabeled word and then determine the word sentiment polarities based on its scores. We rank the two sets of newly labeled sentiment words according to their polarity scores. The top-ranked Chinese words are shown in Table 1. We manually label the top-ranked 1K sentiment words. For P@10K, we sequentially divide the top 10K ranked list into ten equal parts. One hundred sentiment words are randomly selected from each part for labeling. Similar to the evaluation of TREC Blog Distillation (Ounis, Macdonald, and Soboroff 2008), all the labeled words from each approach are used in the evaluation. We then evaluate the ranked lists with two metrics, Precision@K and Recall.

⁶ http://www.globalwordnet.org/gwa/wordnet_table.html.

Table 1
The top learned Chinese sentiment words.

Word	Meaning	Polarity	Word	Meaning	Polarity
好	good	positive	灾难	disaster	negative
正确	correct	positive	悲剧	tragedy	negative
有用	useful	positive	危险	dangerous	negative
聪明	smart	positive	伤害	harm	negative
高兴	happy	positive	错误	fault	negative
可靠	reliable	positive	愤怒	rage	negative
准确	accurate	positive	失败	fail	negative
快乐	happy	positive	破坏	damage	negative
乐观	optimistic	positive	痛苦	sore	negative
忠诚	loyal	positive	冲突	clash	negative

4.2 Evaluation of the Bilingual Word Graph

In this set of experiments, we examine the influence of graph topologies on sentiment lexicon learning.

Mono: This approach learns the Chinese sentiment lexicon based only on the Chinese monolingual word graph $\mathcal{G}_T = (\mathbf{X}_T, \mathbf{W}_T \cup \tilde{\mathbf{W}}_T)$. Because it needs labeled sentiment words, we incorporate the English labeled sentiment words \mathbf{X}_E and the inter-language relation \mathbf{W}_A in the first iteration. Then we set \mathbf{X}_E and \mathbf{W}_A to be zero in later iterations.

BLP-WOA (bilingual word graph without antonym): This approach is based on the bilingual word graph. It only involves the inter-language relation \mathbf{W}_A and the synonym intra-language relations \mathbf{W}_E and \mathbf{W}_T . $\tilde{\mathbf{W}}_E$ and $\tilde{\mathbf{W}}_T$ are set to be zero.

BLP: This approach is based on the bilingual word graph. It incorporates the inter-language relation \mathbf{W}_A , the synonym intra-language relations \mathbf{W}_E and \mathbf{W}_T , and the antonym intra-language relations $\tilde{\mathbf{W}}_E$ and $\tilde{\mathbf{W}}_T$.

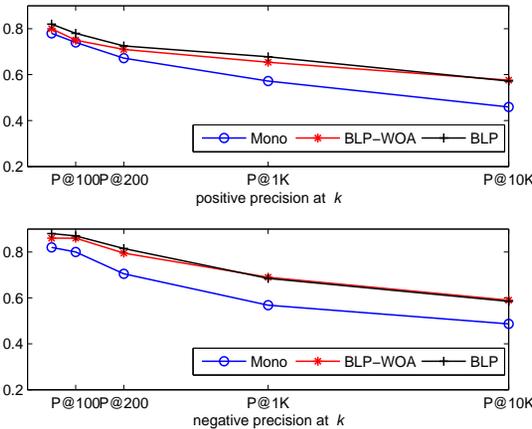


Figure 3
Precision evaluation of the experiments based on monolingual (Mono), bilingual without antonym (BLP-WOA), and bilingual (BLP) word graphs.

Table 2

Recall evaluation of the experiments based on monolingual (Mono), bilingual without antonym (BLP-WOA), and bilingual (BLP) word graphs.

Chinese	<i>Mono</i>	<i>BLP-WOA</i>	<i>BLP</i>
Positive	0.623	0.702	0.708
Negative	0.631	0.706	0.709

In these approaches, μ is set to 0.1 as in Zhou et al. (2004). The precision of these approaches are shown in Figure 3. The figure shows that the approaches based on the bilingual word graph significantly outperform the one based on the monolingual word graph. The bilingual word graph can bring in more word relations and accelerate the sentiment propagation. Besides, in the bilingual word graph, the English sentiment seed words can continually provide accurate sentiment information. Thus we observe the increase in the approaches based on the bilingual word graph in term of both precision and recall (Table 2). Meanwhile, we find that adding the antonym relation in the bilingual word graph slightly enhances precision in top-ranked words and similar findings are observed in our later experiments. It appears that the antonym relations depict word relations in a more accurate way and can refine the word sentiment scores more precisely. However, the synonym relation and word alignment relation dominate, whereas the antonym relation accounts for only a small percentage of the graph. It is hard for the antonym relation to introduce new relations into the graph and thus it cannot help to further improve recall.

4.3 Comparison with Baseline and Existing Approaches

In this set of experiments, we compare our approach with the baseline and existing approaches.

Rule: For the intra-language relation, this approach assumes that the synonyms of a positive (negative) word are always positive (negative), and the antonyms of a positive (negative) word are always negative (positive). For the inter-language relation, we regard the Chinese word aligned to positive (negative) English words as positive (negative). If a word connects to both positive and negative English words, we regard it as objective. Based on this heuristic, we generate two sets of sentiment words.

SOP: Hassan et al. (2011) present a method to predict the semantic orientation of unlabeled words based on the mean hitting time to the two sets of sentiment seed words. Given the graph $\mathcal{G} = (\mathbf{X}_E \cup \mathbf{X}_T, \mathbf{W}_E \cup \tilde{\mathbf{W}}_E \cup \mathbf{W}_T \cup \tilde{\mathbf{W}}_T \cup \mathbf{W}_A)$, it defines the transition probability from node i to node j as

$$p(j|i) = \frac{w_{i,j}}{\sum_k w_{i,k}}$$

The mean hitting time $h(i|j)$ is the average number of the weighted steps from word i to word j . Starting with the word i and ending with the sentiment word $k \in M$, the mean hitting time $h(i|M)$ can be formally defined as

$$h(i|M) = \begin{cases} 0, & i \in M \\ \sum_{j \in V} p(j|i) \times h(j|M) + 1, & \text{otherwise} \end{cases}$$

Let M_+ and M_- denote the GI positive and negative seeds. If $h(i|M_+)$ is greater than $h(i|M_-)$, the word x_i is classified as negative; otherwise it is classified as positive. The generated positive words and negative words are then ranked according to their polarity scores, respectively.

MAD: Talukdar and Crammer (2009) propose a MAD algorithm to modify the adsorption algorithm (Baluja et al. 2008) by adding a new regularization term. In particular, besides the positive and negative labels, a dummy label is assigned to each word in the *MAD* approach. Two additional columns, representing the scores of the dummy label, are added into \mathbf{Y} and \mathbf{F} , respectively. We denote these two matrices with the dummy labels as $\hat{\mathbf{Y}}$ and $\hat{\mathbf{F}}$. Meanwhile, $\hat{\mathbf{R}}$ is used to represent the initial dummy scores of all the words. For a word x_i , the newly added columns in $\hat{\mathbf{Y}}$ and $\hat{\mathbf{F}}$ are set to zero (i.e., $\hat{y}(i, 3) = \hat{f}(i, 3) = 0$). $\hat{r}(i, 0)$ and $\hat{r}(i, 1)$ are set to zero, and $\hat{r}(i, 3)$ is assigned to one. Then, the predicted label $\hat{\mathbf{F}}_i$ of the word x_i is iteratively obtained by

$$\hat{\mathbf{F}}_i^{(t+1)} = \frac{1}{\hat{\mathbf{M}}_{ii}} (\lambda_1 \hat{\mathbf{Y}}_i + \lambda_2 \sum_j \mathbf{W}_{ij} \hat{\mathbf{F}}_j^{(t)} + \lambda_3 \gamma \hat{\mathbf{R}}_i)$$

$$\hat{\mathbf{M}}_{ii} = (\lambda_1 + \lambda_2 \sum_{j \neq i} \mathbf{W}_{ji} + \lambda_3)$$

$\lambda_{1 \sim 3}$ and γ are used to tune the importance of each iteration term. We set $\lambda_{1 \sim 2}$ to one, λ_3 to 10, and γ to 0.1, which produces reasonably good results. After propagation, $\hat{f}(i, 0)$ and $\hat{f}(i, 1)$ are used to determine the sentiment polarity of the word x_i .

We show recall of the learned Chinese sentiment words in Table 3. Compared with *BLP* and *SOP*, the *Rule* approach learns fewer sentiment words. The coverage of the *Rule* approach is inevitably low because many words in the corpus are aligned to both positive and negative words. For example, in most cases the positive Chinese word 帮助 (helpful) is aligned to the positive English word *helpful*. But sometimes it is aligned (or misaligned) to the negative English words, like *freak*. Under this situation, the word tends to be predicted as objective. In *SOP*, the positive and negative scores are related to the distances of the word to the positive and negative seed words, and the distance is usually coarse-grained to depict the sentiment polarity. For example, the shortest path between the word *good* and the word *bad* in WordNet is only 5 (Kamps et al. 2004). The *Rule* and *SOP* approaches find different sentiment words. We then evaluate the learned Chinese polarity word lists by precision at k . As illustrated in Figure 4, the significance test indicates that our approach significantly outperforms the *Rule* and *SOP* approaches. The major difference of our approach is that the polarity information can be transferred between English and Chinese and within each language at the same time, whereas in the other two approaches the polarity information mainly transfers from English to Chinese and once a word gets a polarity score, it is difficult to change or refine. The idea of the *MAD* approach is similar to bilingual graph label propagation, but the *MAD*

Table 3
Recall evaluation of the *Rule*, *SOP*, *MAD*, and *BLP* approaches.

Chinese	<i>Rule</i>	<i>SOP</i>	<i>MAD</i>	<i>BLP</i>
Positive	0.382	0.604	0.662	0.708
Negative	0.371	0.582	0.681	0.709

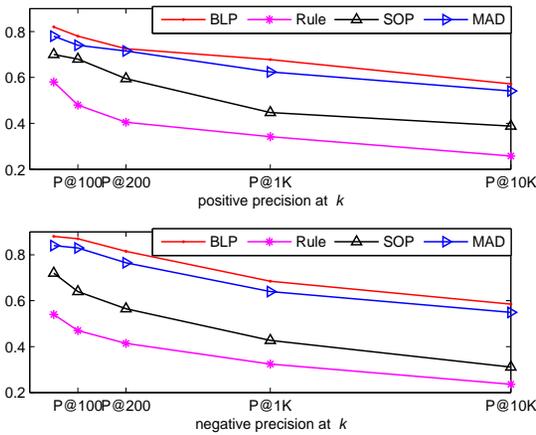


Figure 4
Precision evaluation of the *Rule*, *SOP*, *MAD*, and *BLP* approaches.

approach fails to leverage the antonym intra-language relation. We observe that the *MAD* approach can achieve a comparable result to the *BLP* approach. *MAD* can obtain a smoother label score by adding a dummy label. But the dummy label does not influence the sentiment polarity labels too much because it is not used in the determination of the word sentiment polarity. Besides, *MAD* cannot deal with the antonym relation. As a result, these experiments demonstrate the overall superiority of our approach in cross-lingual sentiment lexicon learning. This also indicates the effectiveness of the *BLP* approach in Chinese sentiment lexicon learning.

4.4 Evaluation of the Inter-Language Relation

This set of experiments is to examine the ways to build the inter-language relation.

BLP-dict: The inter-language relation is built upon the translation entries from LDC⁷ and Universal Dictionary (UD).⁸ From these dictionaries (both English–Chinese and Chinese–English dictionaries), we collect 41,034 translation entries between the English and Chinese words. If the English word x_i can be translated to the Chinese word x_j in UD dictionary, $w_A(i, j)$ and $w_A(j, i)$ are set to 1.

BLP-MT: All the Chinese (English) words are translated into English (Chinese) by Google Translator. If the Chinese word x_i can be translated to the English word x_j , the $w_A(i, j)$ and $w_A(j, i)$ are set to 1. If a Chinese word is translated to an English phrase, we assume that the Chinese word is projected to each word in the English phrase. To improve the coverage, we translate the English sentiment seed words with three other methods; they are word collocation, coordinated phrase, and punctuation, as mentioned in Meng et al. (2012b).

The learned Chinese sentiment word lists are also evaluated with precision at k . As shown in Figure 5, we find that the alignment-based approach outperforms the dictionary-based and MT-based approaches. The reason that contributes to this is that we can build more inter-language relations based on word alignment, compared

⁷ http://projects.ldc.upenn.edu/Chinese/LDC_ch.htm.

⁸ <http://www.dicta.info/uddl.php>.

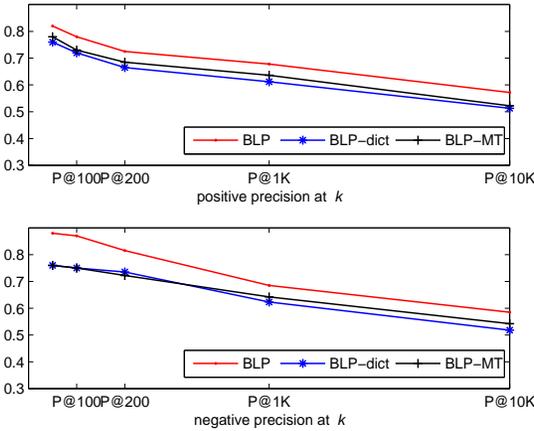


Figure 5
Influence on Precision of the inter-language relation.

with the translation entries from the dictionary and the translation pairs from Google Translator. For example, the English word *move* is often translated to 移动 (shift) and 感动 (affect, touch) by dictionaries or MT engines. From the parallel sentences, besides these word translation pairs, the word *move* can be also aligned to 一帆风顺 (plain sailing bon voyage) that is commonly used in Chinese greeting texts. This translation entry is hard to find in dictionaries or by MT engines. The words are aligned between the two parallel sentences. Sometimes the word *move* may be forced to be aligned to 一帆风顺 in the parallel sentences *good luck and best wishes on your career move* and *祝你|新的|事业|一帆风顺*. Thus, when building the inter-language relations with word alignment, our approach is likely to learn more sentiment word candidates. It is also the reason why the dictionary-based and MT-based approaches learn fewer sentiment words than our approach, as indicated in Table 4. According to our statistic, on average a Chinese word is connected to 2.3 and 2.1 English words if we build the inter-language relations with the dictionary and Google Translator, respectively. By building the inter-language relation with word alignment, our approach connects a Chinese word to 16.21 English words an average, which greatly increases the coverage of the learned sentiment lexicon.

4.5 Evaluation of the Intra-Language Relation

The following set of experiments reveals the influence of the intra-language relation.

BLP-A: As the baseline of this set of experiments, it does not build the intra-language relations with either English or Chinese WordNet synsets. Only the

Table 4
Influence on Recall of the inter-language relation.

Chinese	<i>BLP-dict</i>	<i>BLP-MT</i>
Positive	0.649	0.654
Negative	0.660	0.679

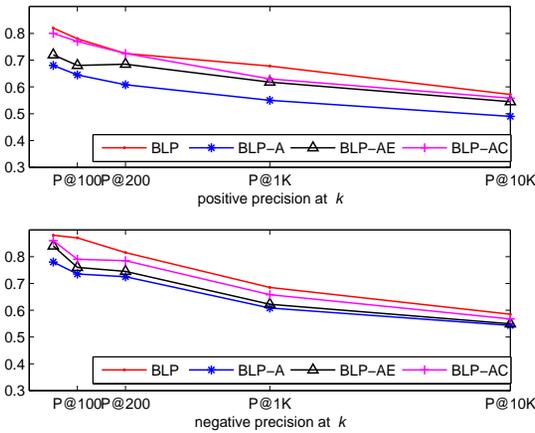


Figure 6
Influence on Precision of the intra-language relation.

inter-language relation with word alignment is used to build the graph. That means \mathbf{W}_E , $\tilde{\mathbf{W}}_E$, \mathbf{W}_T , and $\tilde{\mathbf{W}}_T$ are defined as zero matrixes.

BLP-AE: Word alignment and the English WordNet synsets are used to build the intra-English relation \mathbf{W}_E , but the intra-Chinese relation \mathbf{W}_T and $\tilde{\mathbf{W}}_T$ are set to zero matrixes.

BLP-AC: Word alignment and the Chinese WordNet synsets are used to build the intra-Chinese relation \mathbf{W}_T , but the intra-English relation \mathbf{W}_E and $\tilde{\mathbf{W}}_E$ are set to zero matrixes.

As Figure 6 shows, when combining both English and Chinese intra-language relations, the precision curves of both positive and negative predictions increase. This indicates that adding the intra-language relations has a positive influence. The improvement can be explained by the ability of the intra-language relations to refine the polarity scores. For example, the English word *sophisticated* can be aligned to the positive Chinese word 精致的 (delicate) as well as the negative Chinese word 圆滑的 (wily, wicked). In the GI lexicon, the English word *sophisticated* is labeled as positive. In the bilingual word graph that contains only the inter-language relations, the negative Chinese word 圆滑的 is likely to be labeled as positive. However, with the intra-language relation, the negative Chinese word 圆滑的 may connect to the other negative Chinese words, like 狡猾的 (foxy); and the Chinese positive word 精致的 may connect to the other positive Chinese words, like 精巧的 (elaborate). Thus the polarity score of the word can be refined by the intra-language relation in each iteration of propagation. Another advantage of the intra-language relation is that it helps to reduce the noise introduced by the inter-language relation. For example, sometimes the Chinese positive word 帮助 (help) is misaligned to the negative English word *freak* by the inter-language relation, but it is also connected to the synonyms 有助 (help) and 有益 (salutary) (which are positive) by the intra-language relations. The polarity score of the word 帮助 can be adjusted by the intra-language relation. Thus, though the inter-language relation brings in certain noisy alignments, the intra-language relation can help to refine the polarity score of the word using its intra-language relation.

Table 5
Numbers of labeled sentiment sentences.

	Positive	Negative	Neutral
sentence number	1,218	944	528

4.6 Sensitivity of Parameter $\rho_{1\sim 4}$

ρ_1 and ρ_2 in Equation (3) tune the English and Chinese synonym intra-language propagation, while ρ_3 and ρ_4 in Equation (4) adjust the English and Chinese antonym intra-language propagation. For simplicity, let ρ_1 equal ρ_2 and let ρ_3 equal ρ_4 . Then we tune $\rho_{1,2}$ and $\rho_{3,4}$ together. When $\rho_{1,2}$ and $\rho_{3,4}$ range from $\{1e-2, 1e-1, 1, 10, 100, 1000\}$, Precision@1K ranges from 0.631 to 0.689 and Recall ranges from 0.651 to 0.729 on average. In general, we find that when $1 \leq \rho_{3,4} < \rho_{1,2} \leq 10$, we can obtain better results.

4.7 Evaluation on Sentiment Classification

Sentiment classification is one of the most extensively studied tasks in the community of sentiment analysis (Pang and Lee 2008). To see whether the performance improvement in lexicon learning also improves the results of sentiment classification, we apply the generated Chinese sentiment lexicons to sentence-level sentiment classification.

Data set: The NTCIR sentiment-labeled corpus is used for sentiment classification (Seki et al. 2008, 2009). We extract the Chinese sentences that have positive, negative, or neutral labels. The numbers of extracted sentences are shown in Table 5. The learned sentiment words in the Mono and BLP approaches are used as classification features. We implement the following baselines for comparison.

BSL_{DF}: The Chinese word unigrams and bigrams are extracted from the NTCIR data set as features. We rank the features according to their frequencies and gradually increase the value of N for the Top- N classification features.

BSL_{LF}: The words in existing Chinese sentiment lexicons are used as features. A total of 836 positive words and 1,254 negative words are collected from *HowNet*.⁹

We use *LibSVM*¹⁰ and perform 10-fold cross-validation on the NTCIR polarity sentences. The accuracies over N number of features are plotted in Figure 7. Our approach achieves a very promising improvement, although the features and the sentences that need to be classified are selected from different corpora. This suggests that the generated sentiment lexicon is adaptive and qualitative enough for sentiment classification.

5. Conclusions and Future Work

In this article, we studied the task of cross-lingual sentiment lexicon learning. We built a bilingual word graph with the words in two languages and connected them with the inter-language and intra-language relations. We proposed a bilingual word graph label propagation approach to transduce the sentiment information from English sentiment words to the words in the target language. The synonym and antonym relations among

⁹ <http://www.keenage.com/>.

¹⁰ <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>.

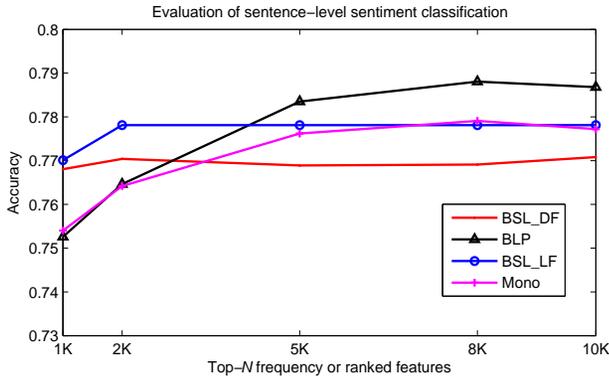


Figure 7
Accuracy evaluation on sentiment classification.

the words in the same languages are leveraged to build the intra-language relations. Word alignment derived from a large parallel corpus is used to build the inter-language relations. Experiments on Chinese sentiment lexicon learning demonstrate the effectiveness of the proposed approach. There are three main conclusions from this work. First, the bilingual word graph is suitable for sentiment information transfer and the proposed approach can iteratively improve the precision of the generated sentiment lexicon. Second, building the inter-language relations with the large parallel corpus can significantly improve the coverage. Third, by incorporating the antonym relations into the bilingual word graph, the BLP approach can achieve an improvement in precision. In the future, we will explore the opportunity of expanding or generating the sentiment lexicons for multiple languages by bootstrapping.

Acknowledgments

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

References

- Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the 7th International Conference on Language Resources and Evaluation*, pages 2,200–2,204, Malta.
- Baluja, Shumeet, Rohan Seth, D. Sivakumar, Yushi Jing, Jay Yagnik, Shankar Kumar, Deepak Ravichandran, and Mohamed Aly. 2008. Video suggestion and discovery for Youtube: Taking random walks through the view graph. In *Proceedings of the 17th International Conference on the World Wide Web*, pages 895–904, Beijing.
- Banea, Carmen, Rada Mihalcea, and Janyce Wiebe. 2010. Multilingual subjectivity: Are more languages better? In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 28–36, Beijing.
- Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Boyd-Graber, Jordan and Philip Resnik. 2010. Holistic sentiment analysis across languages: Multilingual supervised latent Dirichlet allocation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 45–55, Boston, MA.
- Das, Dipanjan and Slav Petrov. 2011. Unsupervised part-of-speech tagging with bilingual graph-based projections. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 600–609, Portland, OR.
- Duh, Kevin, Akinori Fujino, and Masaaki Nagata. 2011. Is machine translation ripe

- for cross-lingual sentiment classification? In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 429–433, Portland, OR.
- Esuli, Andrea and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation*, pages 417–422, Las Palmas.
- Esuli, Andrea and Fabrizio Sebastiani. 2007. Random-walk models of term semantics: An application to opinion-related properties. In *Proceedings of the 3rd Language and Technology Conference*, pages 221–225, Poznan.
- Hassan, Ahmed, Amjad Abu-Jbara, Rahul Jha, and Dragomir Radev. 2011. Identifying the semantic orientation of foreign words. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 592–597, Portland, OR.
- Hatzivassiloglou, Vasileios and Kathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 174–181, Madrid.
- He, Yulan, Harith Alani, and Deyu Zhou. 2010. Exploring English lexicon knowledge for Chinese sentiment analysis. In *Proceedings of the 2010 CIPS-SIGHAN Joint Conference on Chinese Language Processing*, Beijing.
- Hu, Minqing and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177, Seattle, WA.
- Hu, Minqing and Bing Liu. 2006. Opinion extraction and summarization on the Web. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 1,621–1,624, Boston, MA.
- Jiang, Jonathan Q. 2011. Learning protein functions from bi-relational graph of proteins and function annotations. *Algorithms in Bioinformatics*, 6833:128–138, Springer Berlin Heidelberg.
- Jiang, Jonathan Q. and Lisa J. McQuay. 2012. Predicting protein function by multi-label correlated semi-supervised learning. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 9(4):1059–1069.
- Kamps, Jaap, Maarten Marx, Robert J. Mokken, and Maarten de Rijke. 2004. Using Wordnet to measure semantic orientations of adjectives. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, pages 1,115–1,118, Lisbon.
- Kim, Soo-Min and Eduard Hovy. 2004. Determining the sentiment of opinions. In *Proceedings of the 20th International Conference on Computational Linguistics*, pages 355–363, Geneva.
- Kunegis, Jerome, Stephan Schmidt, Andreas Lommatzsch, Jürgen Lerner, Ernesto W. De, and Luca Sahin Albayrak. 2010. Spectral analysis of signed graphs for clustering, prediction and visualization. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, pages 559–570, Columbus, OH.
- Li, Shen, Joao V. Graca, and Ben Taskar. 2012. Wiki-ly supervised part-of-speech tagging. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1,389–1,398, Jeju Island.
- Liang, Percy, Ben Taskar, and Dan Klein. 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*, pages 104–111, New York, NY.
- Lu, Bin, Chenhao Tan, Claire Cardie, and Benjamin K. Tsou. 2011. Joint bilingual sentiment classification with unlabeled parallel corpora. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 320–330, Portland, OR.
- Meng, Xinfan, Furu Wei, Xiaohua Liu, Ming Zhou, Ge Xu, and Houfeng Wang. 2012a. Cross-lingual mixture model for sentiment classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers*, pages 572–581, Jeju Island.
- Meng, Xinfan, Furu Wei, Ge Xu, Longkai Zhang, Xiaohua Liu, Ming Zhou, and Houfeng Wang. 2012b. Lost in translations? Building sentiment lexicons using context-based machine translation. In *the 24th International Conference on Computational Linguistics*, pages 829–838, Bombay.
- Mihalcea, Rada, Carmen Banea, and Janyce Wiebe. 2007. Learning multilingual subjective language via cross-lingual projections. In *Proceedings of the 45th Annual Meeting of the Association for*

- Computational Linguistics*, pages 976–983, Prague.
- Miller, George A. 1995. Wordnet: A lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Munteanu, Dragos Stefan and Daniel Marcu. 2005. Improving machine translation performance by exploiting non-parallel corpora. *Journal of Computational Linguistics*, 31(4):477–504.
- Ounis, Iadh, Craig Macdonald, and Ian Soboroff. 2008. Overview of the TREC-2010 blog track. In *Proceedings of the NTCIR07 workshop*, pages 104–111, Tokyo.
- Pang, Bo and Lillian Lee. 2008. *Opinion mining and sentiment analysis*, volume 2. Foundations and Trends in Information Retrieval. Now Publishers, Inc.
- Qiu, Guang, Bing Liu, Jiajun Bu, and Chun Chen. 2011. Opinion word expansion and target extraction through double propagation. *Computational Linguistics*, 37(1):9–27.
- Rao, Delip and Deepak Ravichandran. 2009. Semi-supervised polarity lexicon induction. In *Proceedings of the 12th Conference of the European Chapter of the ACL*, pages 675–682, Athens.
- Riloff, Ellen, Janyce Wiebe, and Theresa Wilson. 2003. Learning subjective nouns using extraction pattern bootstrapping. In *Proceedings of the 7th Conference on Natural Language Learning*, pages 25–32, Edmonton.
- Saad, Yousef. 2003. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition.
- Seki, Yohei, David Kirk Evans, Lun-Wei Ku, Le Sun, Hsin-Hsi Chen, and Noriko Kando. 2008. Overview of multilingual opinion analysis task at NTCIR-7. *Proceedings of the NTCIR07 Workshop*, pages 104–111, Tokyo.
- Seki, Yohei, Lun-Wei Ku, Le Sun, Hsin-Hsi Chen, and Noriko Kando. 2009. Overview of multilingual opinion analysis task at NTCIR-8: A step toward cross lingual opinion analysis. In *Proceedings of the NTCIR08 Workshop*, pages 209–220, Tokyo.
- Steinberger, Josef, Polina Lenkova, Mohamed Ebrahim, Maud Ehrmann, Ali Hurriyetoglu, Mijail Kabadjov, Ralf Steinberger, Hristo Tanev, Vanni Zavarella, and Silvia Vazquez. 2011. Creating sentiment dictionaries via triangulation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 28–36, Beijing.
- Stone, Philip J. 1997. Thematic text analysis: New agendas for analyzing text content. In *Text Analysis for the Social Sciences*, chapter 2. Lawrence Erlbaum, Mahwah, NJ.
- Takamura, Hiroya, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pages 130–140, Ann Arbor, MI.
- Talukdar, Partha Pratim and Koby Crammer. 2009. New regularized algorithms for transductive learning. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 442–457, Bled.
- Turney, Peter D. and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):315–346.
- Wan, Xiaojun. 2009. Co-training for cross-lingual sentiment classification. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language*, pages 235–243, Singapore.
- Wang, Hua, Heng Huang, and Chris Ding. 2011. Image annotation using bi-relational graph of images and semantic labels. In *the 24th IEEE Conference on Computer Vision and Pattern Recognition*, pages 793–800, Colorado Springs, CO.
- Yu, Hong and Vasileios Hatzivassiloglou. 2003. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 129–136, Seattle, WA.
- Zhou, Dengyong, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Scholkopf. 2004. Learning with local and global consistency. In *Proceedings of Advances in Neural Information Processing Systems*, pages 321–328, Vancouver.
- Zhu, Xiaojin and Zoubin Ghahramani. 2002. Learning from labeled and unlabeled data with label propagation. Technical report CMU-CALD-02-107, Carnegie Mellon University.
- Zhu, Xiaojin, Zoubin Ghahramani, and John Lafferty. 2003. Semi-supervised learning using Gaussian fields and harmonic functions. In *International Conference Machine Learning*, pages 912–919, Washington, DC.