Knowledge-Based Systems 71 (2014) 61–71

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

Data-driven integration of multiple sentiment dictionaries for lexicon-based sentiment classification of product reviews

Heeryon Cho [a], Songkuk Kim [a,b,*], Jongseo Lee [a,b], Jong-Seok Lee [a,b,*]

[a] Yonsei Institute of Convergence Technology, Yonsei University, Republic of Korea
[b] School of Integrated Technology, Yonsei University, Republic of Korea

A R T I C L E   I N F O

Article history:
Available online 13 June 2014

Keywords:
Sentiment dictionary
Domain adaptation
Integration
Lexicon-based review classification
Sentiment analysis

A B S T R A C T

In lexicon-based sentiment classification, the problem of contextual polarity must be explicitly handled since it is a major cause for classification error. One way to handle contextual polarity is to revise the prior polarity of the sentiment dictionary to fit the domain. This paper presents a data-driven method of adapting sentiment dictionaries to diverse domains. Our method first merges multiple sentiment dictionaries at the entry word level to expand the dictionary. Then, leveraging the ratio of the positive/negative training data, it selectively removes the entry words that do not contribute to the classification. Finally, it selectively switches the sentiment polarity of the entry words to adapt to the domain. In essence, our method compares the positive/negative review's dictionary word occurrence ratios with the positive/negative review ratio itself to determine which entry words to be removed and which entry words' sentiment polarity to be switched. We show that the integrated sentiment dictionary constructed using our 'merge', 'remove', and 'switch' operations robustly outperforms individual dictionaries in the sentiment classification of product reviews across different domains such as smartphones, movies, and books.

1. Introduction

Reading online product reviews is now a routine part of purchasing experience that provides prospective buyers valuable information regarding the product. Product reviews can be divided into positive and negative reviews in which the reviewer either recommends or does not recommend a product. Due to the ever increasing volume of online product reviews, efficiently dividing these reviews into positive and negative reviews is a fundamental step in utilizing product reviews. A growing body of research on automatically classifying the sentiment in product reviews exists in the field of "*opinion mining and sentiment analysis*" with different solutions to supervised, unsupervised, semi-supervised, and concept-based approaches [1–4]. This study presents yet another lexicon-based approach to classifying product reviews using existing sentiment lexicons, but proposes a novel method of merging and revising multiple sentiment lexicons by incorporating labeled product reviews to enable domain adaptation of sentiment values.

Numerous sentiment lexicons with varying format and size have been constructed to aid the classification of positive and negative sentiments in texts. Examples include SentiSense [5], SentiWordNet [6], Micro-WNOp [7], WordNet-Affect [8], which are based on the English lexical database WordNet [9], SO-CAL [10], AFINN [11], Opinion Lexicon [12], Subjectivity Lexicon [13], General Inquirer [14], which are manually or semi-automatically constructed, and SenticNet [15] that is constructed from the Open Mind common sense knowledge base.

With diverse sentiment lexicons readily available, we are naturally inclined to ask the following questions: (1) How are they different? (2) Can we construct a better sentiment lexicon for sentiment classification by integrating multiple sentiment lexicons? We answer these questions by first comparing different sentiment lexicons and their positive/negative product review classification performances across different domains. We then present *merge*, *remove*, and *switch* operations that merge and revise the entry words of the multiple sentiment lexicons using labeled product reviews. Consequently, an integrated sentiment dictionary is constructed. We show that the integrated sentiment dictionary outperforms individual dictionaries in the sentiment classification tasks of the three kinds of product reviews robustly via domain adaptation.

* Corresponding authors. Address: 162-1 Songdo-dong, Yeonsu-gu, 406-840 Incheon, Republic of Korea. Tel.: +82 32 749 5842 (S. Kim). Tel.: +82 32 749 5846; fax: +82 32 818 5801 (J.-S. Lee).

E-mail addresses: songkuk@yonsei.ac.kr (S. Kim), jong-seok.lee@yonsei.ac.kr (J.-S. Lee).

http://dx.doi.org/10.1016/j.knosys.2014.06.001
0950-7051/

While many existing works approach the acquisition of domain-specific sentiment lexicon as a lexicon *building* problem, the proposed approach starts out with the existing sentiment lexicons, i.e., the final product of the lexicon building approaches, and combines these existing lexicons to come up with improved sentiment lexicons. The main contribution of this study is the development of a novel method of removing and switching the content of the existing sentiment lexicons. While the sentiment lexicons provide a list of sentiment words, the sentiment values are defined a priori as general-purpose values without consideration to a specific domain. Therefore, it is necessary to adjust the sentiment lexicons to the target domain before utilizing them for sentiment classification. Our method first selectively removes the sentiment words from the existing lexicon to prevent erroneous matching of the sentiment words during lexicon-based sentiment classification. Next, it selectively switches the polarity of the sentiment words to adjust the sentiment values to a specific domain. The *remove* and *switch* operations are performed using the target domain's labeled data, online product reviews in this study, by comparing the positive and negative distribution of the labeled reviews with a positive and negative distribution of the sentiment words. We propose a data-driven approach to automatically determine the necessary parameter for removing less contributing sentiment words and switching the polarity of the selected sentiment words to update the sentiment lexicon to fit the given domain. The threshold for judging the positivity or the negativity of given product reviews is also set automatically. The data-driven nature of the proposed approach is feasible and beneficial in situations where the labeled data is abundantly available but a human expert is absent or costly. Although the creation of a large amount of labeled data may be considered as laborious in general, the explosion of labeled Web opinion data and openly available lexical resources (e.g., sentiment lexicons) and the progress in big data processing techniques in recent years have formed an advantageous environment for the data-driven sentiment classification. We fully leverage such an environment in this study. We also conduct extensive experiments to compare the performances of different state-of-the-art sentiment dictionaries, examine the effectiveness of each of the *remove* and *switch* operations when revising the dictionary, analyze the effect of different thresholds and slider values (i.e., parameters used in sentiment classification and revising the dictionary, respectively), and examine the revised dictionaries' performance with respect to the size of the training data.

The remainder of the paper is structured as follows. We outline the existing approaches on sentiment lexicon construction, domain adaptation, and sentiment classification in Section 2, and describe the ten sentiment lexicons and the procedures for standardizing, merging, and revising multiple lexicons in Section 3. We then summarize the experimental setup for sentiment dictionary evaluation in Section 4 and report the sentiment classification performances of the ten individual sentiment dictionaries and one integrated dictionary on three product review classification tasks in Section 5. Finally, we conclude this paper in Section 6.

2. Related work

One of the primary lexicon-based approaches to classifying sentiments in texts focuses on detecting the polarity of a given word using a domain corpus. For instance, Turney and Littman proposed a method of inferring the sentiment of a given word using pointwise mutual information (PMI) by associating the word with a set of positive and negative seed words [16]. The sentiment of a word differs according to different domains, however, and this has led to an extensive research on domain-specific sentiment lexicon extraction and construction. Fahrni and Klenner proposed a domain-specific adaptation of sentiment-bearing adjectives [17]. Adjectives such as "warm" and "cold" possess prior polarity, but depending on the context, this polarity may change. For example, warm mittens may be desirable, but warm beer may not be. To tackle such a problem of contextual polarity, Fahrni and Klenner implemented a two-stage process that first identifies domain-specific targets using Wikipedia, and then determines target-specific polarity of adjectives using a tagged corpus. Unlike this method, our method does not limit the scope of sentiment analysis only to adjectives but can consider contextual polarity of nouns, verbs, and adverbs in addition to adjectives covered by the default sentiment dictionary, which would lead to more accurate analysis. Yu et al. measured the similarity between the seed words and the words in a news corpus by comparing their contextual distribution using an entropy measure to extract useful emotion words [18]. Whereas Yu et al. employed a human annotator to select seed words to ensure the quality of the seed words, our method does not require any human intervention, but allows a fully automatic identification of domain-relevant and domain-irrelevant emotion words. Huang et al. used generic sentiment lexicon, "and" and "but" clues, and synonym and antonym relations to extract pairwise candidate sentiment terms and propagated this knowledge to other sentiment terms to obtain a domain-specific sentiment lexicon [19]. No lexical clues, but only distributional information of positive/negative reviews and review words is employed in our approach, which enables a much simpler processing of sentences. Qiu et al. proposed an approach to iteratively extract both the domain sentiment words and features [20]. In doing so, they parsed dependency relation and part-of-speech information in order to map those types of information to the rules for sentiment word and feature extraction. We, on the other hand, parse only the part-of-speech information for looking up the matching entry words in the sentiment dictionary; our approach performs a simpler and faster processing of sentences. Kanayama and Nasukawa proposed an unsupervised lexicon building method that uses context coherency, i.e., the tendency for the same polarities to appear successively in contexts, in order to detect polar clauses that convey positive or negative aspects in a specific domain [21]. They treat Japanese polar clauses as unit of sentiment analysis and use domain-independent polar clauses and conjunctions such as "and", "but", and "because" as clues for polarity detection. While their method is also fully automatic and leverages the distribution of the context coherency, their method also requires syntactic parsing of the sentences. Neviarouskaya et al. proposed methods for automatically expanding a sentiment dictionary by adding and scoring new words based on the sentiment-scored lemmas and types of affixes and by leveraging the direct synonym/antonym/hyponym relations, derivation, and compounding with known lexical units [22]. Their approach, however, focuses on the expansion of the prior-polarity sentiment words with no special concerns to domain-specific polarity. Our approach, on the other hand, considers both the domain-specific polarity and the expansion of the sentiment dictionary by merging and revising multiple sentiment dictionaries. Note that in the process, we generate an improved sentiment dictionary by integrating and revising multiple sentiment dictionaries although our goal is not about building a sentiment dictionary per se. Lu et al. proposed an optimization framework that provides a unified and principled way to combine general-purpose sentiment lexicon, overall sentiment rating, thesaurus, and linguistic heuristics for learning context-dependent sentiment lexicon [23]. We do not employ any lexical heuristics, but simply leverages statistical information, i.e., the positive and negative ratio of the product reviews and the review-matched dictionary entry words. Most of the approaches mentioned above focus on building and expanding the domain-specific sentiment lexicon using combinations of domain corpus,