

TSA'09 Workshop Summary: Topic-Sentiment Analysis

Bei Yu
School of Information Studies
Syracuse University
320 Hinds Hall
Syracuse, NY 13244
byu@syr.edu

Maojin Jiang
Computer Science Department
Illinois Institute of Technology
10 West 31 St, Stuart Building 235
Chicago, IL 60616
jianmao@iit.edu

ABSTRACT

This workshop seeks to bring together researchers in both computer science and social sciences who are interested in developing and using topic-sentiment analysis methods to measure mass opinion, and to foster communications between the research community and industry practitioners as well.

Categories and Subject Descriptors

I.2.7 [Artificial Intelligence]: Natural Language Processing – *text analysis*; I.7.5 [Document and Text Processing]: Document Capture – *document analysis*.

General Terms

Algorithms, Measurement, Performance, Design, Languages, Theory.

Keywords

Sentiment classification, opinion analysis

1. INTRODUCTION

The increasing amount of user-generated content on the Internet and social media and the digitization of large number of government and institutional documents provide new opinion-rich data sources for researchers to examine individual and group perceptions on products, organizations, and social issues at a large scale, and thus contribute to the research and practice in the areas of political science, social policy, communications, and business intelligence.

On the other hand, researchers are tackling the problem of processing large amount of opinion-rich data using various approaches. The increasing number of relevant publications in top data mining, information retrieval and natural language processing conferences (KDD, SIGIR, ACL, WWW, etc.) has witnessed the growing interest in automatic opinion analysis. Both TREC and TAC (Text Analysis Conference) have set up individual tracks for opinion retrieval and analysis tasks.

In recent years, topic detection and tracking techniques have been well developed to identify the issues discussed in a large text collection. Sentiment analysis is catching up to detect the polarity of opinions expressed in texts. However, many times real-world applications have to take into consideration of both topics and

sentiments for precise opinion measurement. Topic and sentiment alignment is crucial for opinion retrieval, extraction, categorization, and aggregation on various issues. Topics and sentiments could also have sophisticated interactions. For example, the choice of topics and the attention distribution among topics might bear hidden opinions as well.

How do we build synergistic topic and sentiment models for text documents? How do we tackle the domain-dependency problem of sentiment analysis? How do we identify users' needs and integrate them into the design of opinion analysis systems? What are the successful applications of topic-sentiment analysis for mass opinion measurement? What lessons have the pioneers learned? How do we evaluate the automatic mass opinion measuring tools with regard to the reliability and validity? This workshop solicits submissions to address these problems and more.

We hope this workshop can advance research in topic-sentiment analysis, make connections between research community and industry practitioners and encourage development of high performance tools and systems that can work at the web scale for real world applications.

2. WORKSHOP OVERVIEW

The call for papers attracted 21 submissions from 14 countries. The program committee accepted 12 papers that cover a variety of topics, including topic-sentiment modeling, sentiment classification and retrieval, sentiment corpus construction, and applications of topic-sentiment analysis in text summarization, question answering, and recommender systems.

The proposed approaches analyze opinions at all levels of granularity: clause, sub-sentence, sentence, paragraph and document. Most of the approaches combine machine learning and statistical methods and the use of linguistic resources (sentiment lexicons, syntactic rules, etc.) for sentiment identification. User-generated content (UGC) is still the main source of data for topic-sentiment analysis in various domains, like customer reviews, blogs, and discussion boards. In addition to English, corpora of other languages (Chinese, Spanish, and Portuguese) have also been studied. Many authors chose to manually annotate their own sentiment corpora to train machine learning algorithms, or employed automatic methods to acquire the sentiment annotation. This indicates the strong demand for large volume of annotated data in various topics and domains to facilitate topic-sentiment analysis.

Figure 1 summarizes the above paragraph by visualizing the keywords of the accepted papers in a word cloud generated by <http://www.tagmap.org>. The three keyword clusters surrounding

“sentiment” and “opinion” indicate the popular sentiment data sources (e.g. “reviews” and “blogs”), the common topic-sentiment analysis tasks (e.g. “classification” and “retrieval”), and the proposed approaches (e.g. “context” and “dependency”) respectively.



Figure 1. Word cloud of the keyword clusters

3. WORKSHOP ORGANIZATION

This workshop consists of one keynote speech and three sessions of paper presentations. Professor Claire Cardie gave the keynote speech. We summarized three main themes among the accepted papers and did our best to assign the papers to the corresponding sessions. Of course some papers cover more than one theme or do not fit in any of these themes. Here we introduce these papers by their main themes and relatedness to other papers.

A large percent of the workshop papers focus on modeling the topic-sentiment interaction (Theme 1). Some papers deal with the main challenges in sentiment analysis, such as domain dependency, complex opinions with multiple aspects and multiple polarities, efficient annotation acquisition, and irony detection (Theme 2). Some papers aim to integrate topic-sentiment analysis as a component to support and improve other text mining systems, such as question answering and recommender systems (Theme 3).

Papers in Theme 1 focus on the interaction between topic and sentiment. Brooke and Hurst explored the interaction of polarity, topic, and discourse in a large video review corpus. Their results indicate important regularities in the expression of opinions and topics, for example concentration of opinions and topic aspects at the beginning and the end of text units, and predictable opinion and topic transition indicated by discourse cues. O’Hare et al. aimed to detect the polarity of financial blogs toward companies and their stocks. The authors developed their own corpus to train sentiment classifier, and proposed text extraction techniques to create topic-specific sub-documents to deal with the topic shift

problem in blog articles. Under similar concept, Zhu et al. took a different approach to segment sentences which contain multiple topic aspects or polarities. Nowson argued that the features used for sentiment classification should be tailored to the topic of a text, and proposed a method to select topic-tuned features to improve sentiment classification accuracy.

Papers in Theme 2 focus on specific challenges in topic-sentiment analysis, such as efficient construction of annotated corpus and other linguistic resources, domain dependency, and irony detection. Sarmento et al. proposed a method which uses a small number of manually-crafted high-precision rules to retrieve opinion-bearing sentences in user comments on news articles, and then quickly increases the size of the corpus by finding more opinionated sentences using opinion propagation patterns. Read and Carroll extended Turney’s methods and examined the usefulness of several forms of word similarity measures for domain-independent sentiment classification. Choi, Kim, and Myaeng tackled the domain dependency problem by training domain-specific sentiment classifiers using contextual features in news texts. Okamoto, Honda, and Eguchi also utilized contextual information to smooth their language models to retrieve sentiment sentences. Although the number of ironic expressions is usually small for most of the sentiment corpora, irony detection is one of the toughest tasks in sentiment analysis. Carvalho et al. suggested that simple gestural patterns are more effective than deeper linguistic information in detecting ironic sentences. On the other hand, deep linguistic analysis was valued by Thet, Na, and Khoo in their approach of clause-level sentiment analysis after splitting the sentences into clauses using dependency trees. There is no doubt that the “text mining” vs. “natural language processing” debate will continue in the topic-sentiment analysis field.

Papers in Theme 3 focus on the integration of topic-sentiment analysis in other text mining applications. Balahur et al. studied the requirements for effectiveness opinion question-answering systems after comparing an open domain QA system and an opinion QA system. Jakob et al. used the opinions extracted from movie reviews to enhance movie recommendation.

4. ACKNOWLEDGMENTS

Putting together TSA 2009 was a team effort. First of all, we would like to thank all authors for their submissions. We would like to express our gratitude to the program committee, who worked very hard in reviewing papers and providing suggestions for their improvements. We would also like to thank the CIKM Workshop Chairs, Min Song and Wook-Shin Han, for their support in the workshop organization process. Special thanks go to Professor Claire Cardie for her keynote speech at our workshop.