Sentiment Analysis of Product Reviews

Cane W. K. Leung
The Hong Kong Polytechnic University, Hong Kong SAR

Stephen C. F. Chan
The Hong Kong Polytechnic University, Hong Kong SAR

INTRODUCTION

Sentiment analysis is a kind of text classification that classifies texts based on the *sentimental orientation* (SO) of opinions they contain. Sentiment analysis of product reviews has recently become very popular in text mining and computational linguistics research. The following example provides an overall idea of the challenge. The sentences below are extracted from a movie review on the Internet Movie Database:

"It is quite boring...... the acting is brilliant, especially Massimo Troisi."

In the example, the author stated that "it" (the movie) is quite boring but the acting is brilliant. Understanding such sentiments involves several tasks. Firstly, evaluative terms expressing opinions must be extracted from the review. Secondly, the SO, or the polarity, of the opinions must be determined. For instance, "boring" and "brilliant" respectively carry a negative and a positive opinion. Thirdly, the opinion strength, or the intensity, of an opinion should also be determined. For instance, both "brilliant" and "good" indicate positive opinions, but "brilliant" obviously implies a stronger preference. Finally, the review is classified with respect to sentiment classes, such as *Positive* and *Negative*, based on the SO of the opinions it contains.

BACKGROUND

Sentiment analysis is also known as opinion mining, opinion extraction and affects analysis in the literature. Further, the terms *sentiment analysis* and *sentiment classification* have sometimes been used interchangeably. It is useful, however, to distinguish between two subtly different concepts. In this article, hence, sentiment analysis is defined as a complete process of extracting and understanding the sentiments being expressed in text documents, whereas sentiment classification is the task of assigning class labels to the documents, or segments of the documents, to indicate their SO.

Sentiment analysis can be conducted at various levels. Word level analysis determines the SO of an opinion word or a phrase (Kamps et al., 2004; Kim and Hovy, 2004; Takamura and Inui, 2007). Sentence level and document level analyses determine the dominant or overall SO of a sentence and a document respectively (Hu and Liu, 2004a; Leung et al., 2008). The main essence of such analyses is that a sentence or a document may contain a mixture of positive and negative opinions. Some existing work involves analysis at different levels. Specifically, the SO of opinion words or phrases can be aggregated to determine the overall SO of a sentence (Hu and Liu, 2004a) or that of a review (Turney, 2002; Dave et al., 2003; Leung et al., 2008).

Most existing sentiment analysis algorithms were designed for binary classification, meaning that they assign opinions or reviews to bipolar classes such as *Positive* or *Negative* (Turney, 2002; Pang et al., 2002; Dave et al., 2003). Some recently proposed algorithms extend binary sentiment classification to classify reviews with respect to multi-point rating scales, a problem known as *rating inference* (Pang and Lee, 2005; Goldberg and Zhu, 2006; Leung et al., 2008). Rating inference can be viewed as a multi-category classification problem, in which the class labels are scalar ratings such as 1 to 5 "stars".

Some sentiment analysis algorithms aim at summarizing the opinions expressed in reviews towards a given product or its features (Hu and Liu, 2004a; Gamon et al., 2005). Note that such *sentiment summarization* also involves the classification of opinions according to their SO as a subtask, and that it is different from classical document summarization, which is about identifying the key sentences in a document to summarize its major ideas.

Sentiment analysis is closely related to *subjectivity analysis* (Wiebe et al., 2001; Esuli and Sebastiani, 2005). Subjectivity analysis determines whether a given text is subjective or objective in nature. It has been addressed using two methods in sentiment analysis algorithms. The first method considers subjectivity analysis a binary classification problem, for example, using *Subjective* and *Objective* as class labels. Pang and Lee (2005) adopted this method to identify subjective sentences in movie reviews. The second method makes use of part-of-speech (POS) information about words to identify opinions (Turney, 2002; Hu and Liu, 2004a; Leung et al., 2008) as previous work on subjectivity analysis suggests that adjectives usually have strong and significant correlation with subjectivity (Bruce and Wiebe, 1999; Wiebe et al., 2001).

MAIN FOCUS

Figure 1 depicts a typical sentiment analysis model. The model takes a collection of reviews as input and processes them using three core steps, *Data Preparation*, *Review Analysis* and *Sentiment Classification*. The

Figure 1. A typical sentiment analysis model

results produced by such a model are the classifications of the reviews, the evaluative sentences, or opinions expressed in the reviews.

Data Preparation

The data preparation step performs necessary data preprocessing and cleaning on the dataset for the subsequent analysis. Some commonly used preprocessing steps include removing non-textual contents, markup tags (for HTML pages) and other information that is not required for sentiment analysis, such as review dates and reviewers' names.

Data preparation may also involve the sampling of reviews for building a classifier. Positive reviews often predominate in review datasets as reported in several studies (e.g. Turney, 2002; Dave et al., 2003; Gamon et al., 2005). Some researchers therefore used review datasets with balanced class distributions when training classifiers to help demonstrate the performance of their algorithms (Pang et al., 2002; Leung et al., 2008).

Review Analysis

The review analysis step analyzes the linguistic features of reviews so that interesting information, including opinions and/or product features, can be identified. This step often applies various computational linguistics tasks to reviews first, and then extracts opinions and product features from the processed reviews. Two commonly adopted tasks for review analysis are POS tagging and negation tagging. POS tagging helps identifying interesting words or phrases having particular POS tags or patterns from reviews (Turney, 2002; Hu and Liu, 2004a; Leung et al., 2008), while negation tagging is used to address the contextual effect of negation words, such as "not", in a sentence (Pang et al., 2002; Dave et al., 2003; Leung et al., 2008). For example, "good" and "not good" obviously indicate opposite SO. Given the term "not good", negation tagging recognizes the existence of the word "not" and adds a special negation tag to the word "good" based on some heuristics.

The review analysis step then proceeds to extract opinions and/or product features from the processed reviews. The opinions or features extracted may be *n*-grams, which are *n* adjacent or nearby words in a sentence (e.g. Turney, 2002). Pang et al. (2002) make use of corpus statistics and human introspection to decide
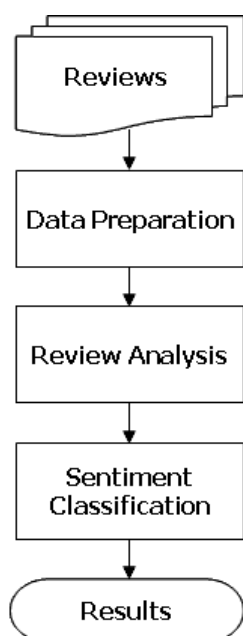
Related Content

New Opportunities in Marketing Data Mining
Victor S.Y. Lo (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition (pp. 1409-1415).*
www.irma-international.org/chapter/new-opportunities-marketing-data-mining/11006

Process Mining to Analyze the Behaviour of Specific Users
Laura Maruster (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition (pp. 1589-1597).*
www.irma-international.org/chapter/process-mining-analyze-behaviour-specific/11031

Program Comprehension through Data Mining
Ioannis N. Kouris (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition (pp. 1603-1609).*
www.irma-international.org/chapter/program-comprehension-through-data-mining/11033

Guide Manifold Alignment by Relative Comparisons
Liang Xiong (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition (pp. 957-963).*
www.irma-international.org/chapter/guide-manifold-alignment-relative-comparisons/10936

Reflecting Reporting Problems and Data Warehousing
Juha Kontio (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition (pp. 1682-1688).*
www.irma-international.org/chapter/reflecting-reporting-problems-data-warehousing/11044