

Designing a Ranking System for Product Search Engine Based on Mining UGC

Dr. Samira Chaabna¹, Professor Wang Hu² & Dr. Mohammed Lutf³

Abstract

With the spread of e-commerce platforms, it becomes extremely difficult for the consumer to choose the right product from a large number of identical products, and different sellers based only on his/her own experience, product pictures or product metadata. Customers reviews present a rich and valuable source of information for potential consumers and manufacturers which can be very helpful in the decision process, but reading all of the available reviews is a hard task and time-consuming. Thus, the automated mining of these reviews and extracting product features in order to generate a ranking system present a valuable and useful tool for consumers to make a well-informed decision. In this paper, we propose a ranking system for product search engine based on mining customers reviews. The product features pairs are extracted using Stanford typed dependencies. The evaluation process of the proposed system has passed through two levels: first measuring of accuracy of product feature extraction and classification, and the second one is the determination of the efficiency of the search results and the usability of visual summarization. The results show a high level of accuracy in feature/opinion pairs extraction and high level of participant satisfaction with the ranking and the summarization.

Keywords: opinion mining, opinion summarization, product search engine, Stanford typed dependencies

1. Introduction

Over the past two decades, the web has seen a tremendous growth and has become one of the major sources of information.

¹ School of Management, Wuhan University of Technology, Wuhan, China. Email: samira.chaabna@gmail.com, Tel: +86-13554307905

² School of Management, Wuhan University of Technology, Wuhan, China. Email: wanghu61@126.com

³ School of Electronic Information and Communications, Huazhong University of Science and Technology, Wuhan, China. Email: mohammed.lutf@gmail.com

In addition, search engines have become the most frequently used tools for searching and accessing information, and a significant channel for promoting and selling products (Ghose, Ipeirotis, and Li, 2013). Online systems that help users to select the most preferential items from a large electronic catalog are known as *product search* and *recommender systems* (Pu, Chen, and Kumar, 2008). These systems aim at decreasing consumer's time and purchase risk by increasing decision accuracy.

There are two main types of information on the web: facts and opinions (Liu, 2006). Currently, most of the information search engines, e.g. Google, search for facts and do not search for opinions. Like the general web search, one can also crawl the user generated content (UGC) on the web and provide an opinion search service as we are going to do in this research. Thus, determining the documents or sentences being used for an opinion search should be done before the search is performed. In our system, the data have been crawled from the leading e-commerce website in China "taobao.com". The objective of an opinion search is to enable users to search for opinions on any object or on any aspect of the object.

Results ranking is crucial to the success of a search engine because users look at the first few results only (Jansen, Spink, and Pedersen, 2005; Silverstein, Marais, Henzinger, and Moricz, 1999). In information search engines (e.g., Google) the ranking of the search results is an immediate signal of the relevance of the result to the query. However, in product search engines, the ranking of the displayed products is often based on criteria such as price, product rating, etc. In our research, the ranking of the proposed system is based on the product features which extracted from the users reviews.

In addition, there are two kinds of opinion search queries: (1) search for opinions on a particular entity or an aspect of the entity, e.g., customer opinions on a cellphone or the battery life of a cellphone, where the query consists of the name of the object and features of the object; (2) the search to all opinions of a person or organization (opinion holder) about a particular entity or an aspect of the entity.

In this type of query, the user may give the name of the opinion holder with the name of the wanted object. The proposed system supports the first type of opinion search queries.

Moreover, with the spread of e-commerce platforms, more customers are turning towards online shopping because it is convenient, fast, and reliable (Zhang, Narayanan, and Choudhary, 2010). Furthermore, the information provided by peer customers is considered more honest, unbiased, comprehensive and describes the usage experience and perspective of (non-expert) customers with similar needs compared to the descriptions provided by the seller (Hu and Liu, 2004b; Wang, Zhu, and Li, 2013).

In order to exploit the feedback provided by customers in the proposed system, it is important to develop an intelligent system that first determine what the feedback is about (which in this research is determining the product features), and then whether this feedback is positive or negative (which in this research is determining the opinion orientation about product features). The first situation is addressed by text mining tools, while the determination of whether the orientation of the feedback is positive or negative is handled by automatic sentiment classification.

In this paper, we address the following two questions: How could the product search engine provide the most appropriate set of products and rank them in a way that satisfies the consumer preferences based on mining customer reviews? What is the best way to extract feature opinion pairs and use them in the ranking of the set of products? In order to answer the above questions, we built a product features-based ranking system for product searching called "Tsearch". The proposed system improves the search result by considering the product features preferred by consumers during the search process and by basing the search result on the previous consumers' opinions about the product instead of focusing only on the information provided by the sellers of the products. Furthermore, the proposed system provides a visual opinion summarization for each product item in order to help the customer gain a general idea about the overall opinions and to determine the most important features that have gained much more attention within a particular product. However, we classify the product features into two categories: product-dependent features (product quality feature) and product-independent features (associated services features), and we have assigned different weights to each category.

Stanford typed dependencies representation was used in order to extract product feature-opinion pairs from customers' reviews. These dependencies were designed to provide a simple description of the grammatical relationships in a sentence that can easily be understood and effectively used by people without linguistic expertise who want to extract textual relations (De Marneffe and Manning, 2008). In this paper we have considered five kinds of dependencies to select candidate product feature/opinion pairs, they are: nsubj, dobj, ccomp, nn, and attr. The evaluation process of the proposed system has passed through two levels: the first one is to measure the accuracy of product feature extraction and classification, and the second one is to determine the efficiency of the search results and the usability of visual summarization. The results show a high level of accuracy in feature opinion pairs extraction and a high level of participants satisfaction with the ranking and the summarization.

The significance of the proposed work compared to previous works in this field of research is summarized in the following aspects. First, this research targets Chinese language, whereas most of the research to date has focused on English and there are no effective resources or tools in other languages that can be used to extract features and determine the sentiment classification effectively. Second, our work considers both product quality features and the associated services features and assigns different weights for each, because the consumers are looking not only for product quality but also for the associated services to satisfy their requirements. In addition, most research to date regarding aspect-based ranking used movie, hotel or restaurant reviews, while the proposed work aims at ranking products based on customer reviews. Finally, providing a visual summarization for each product page is another contribution of this paper, where the user can clearly see the strengths and weaknesses of the product and its aspects in the minds of existing consumers.

The remainder of this paper is organized as follows: related works are presented in the next section. The methodology is presented in section 3. Following that, the proposed approach is introduced in section 4. The detailed implementation of the system is discussed in section 5, followed by the evaluation of the proposed system in section 6. The conclusion and future directions for our work are presented in section 7.

2. Literature Review

With the existence of numerous forums, social networks, discussion groups, and blogs, individual users are participating more actively and are generating a vast amount of new data, termed as *user generated contents* or *customer feedbacks* (Abulaish, Jahiruddin, Doja, and Ahmad, 2009). These feedback has become an important source of information for businesses in developing marketing and product development plans, and for other customers in their process of making a purchase decision.

In recent years, there has been a growing interest in analyzing evaluative texts on the web (Abbasi, Chen, and Salem, 2008; Bafina and Toshniwal, 2013; Pang, Lee, and Vaithyanathan, 2002; Turney, 2002; Wiebe, Wilson, Bruce, Bell, and Martin, 2004). However, the data sources that have been used in literature in performing opinion analysis or product ranking were from blogs (Godbole, Srinivasaiah, and Skiena, 2007; Hui and Gregory, 2010; Melville, Gryn, and Lawrence, 2009; Zhang and Zhou, 2011), review sites (Bafina and Toshniwal, 2013; Becker and Aharonson, 2010), web forums (Abbasi et al., 2008; Hariharan and Ramkumar, 2011; Hariharan, Srimathi, Sivasubramanian, and Pavithra, 2010; Hu, Weng, Zhang, and Xue, 2009), wikis (Mukherjee and Bhattacharyya, 2012), and social networks (Aisopos, Papadakis, Tserpes, and Varvarigou, 2012; Aisopos, Papadakis, and Varvarigou, 2011; Davidov, Tsur, and Rappoport, 2010; Ghiassi, Skinner, and Zimbra, 2013; Jiang, Yu, Zhou, Liu, and Zhao, 2011; Marchetti-Bowick and Chambers, 2012; Montejo-R´aez, Mart´inez-C´amara, Mart´ın-Valdivia, and Uren˜a L´opez, 2014; Pandarachalil and Sendhilkumar, 2013; Souza and Vieira, 2012; Speriosu, Sudan, Upadhyay, and Baldrige, 2011; Tan et al., 2011; Xiang, Fan, Wang, Hong, and Rose, 2012).

In this section, we will cover the main works conducted in the following fields of research: mining and summarizing opinions in general and mining and summarizing Chinese opinions, and finally the literature related to feature-based product ranking.

2.1. Mining and Summarizing Opinions

Opinion mining, also called *sentiment analysis*, is the field of study that analyzes people's opinions, sentiments, attitudes, emotions, evaluations, and appraisals towards entities such as products, services, organizations, individuals, issues, events, topics, and their attributes (Liu, 2012). There are also several names and different tasks considered under the umbrella of sentiment analysis or opinion mining, for example, opinion extraction, sentiment mining, emotion analysis, review mining, subjectivity analysis, affect analysis. In industry, the term "sentiment analysis" is more commonly used, while in academia both "sentiment analysis" and "opinion mining" are frequently employed, but they basically represent the same field of study (Liu, 2011). The term "sentiment analysis" perhaps first appeared in (Nasukawa and Yi, 2003), and the term opinion mining first appeared in (Dave, Lawrence, and Pennock, 2003). However, research on sentiments and opinions appeared earlier (Das and Chen, 2001; Morinaga, Yamanishi, Tateishi, and Fukushima, 2002; Pang et al., 2002; Tong, 2001; Turney, 2002; Wiebe, 2000). In this research, we use the terms sentiment analysis and opinion mining interchangeably.

Sentiment words or *opinion words* are the highest indicators of the sentiments of the opinion holder; some of them are commonly used to express positive or negative sentiments, e.g. good, amazing, bad, terrible. *Sentiment lexicon*, also called *opinion lexicon* consists of a list of such words and phrases. *Sentiment shifters* are expressions that are used to change the sentiment orientations, e.g., from positive to negative or vice versa. *Negation words* are the most important class of sentiment shifters. Sentiment classification attempts to determine whether a text is objective or subjective, which is called *subjectivity classification*, or whether a subjective text contains positive or negative sentiments, which is called *polarity classification (orientation)* (Abbasi et al., 2008; Liu, 2012).

Past researches mainly dealt with three kinds of target units for sentiment classification, i.e. determining whether a word, a sentence or an overall document could be classified as positive or negative.

The task at the document level is to determine whether a whole opinion document (e.g. a review, a blog) is a positive or negative sentiment (Beineke, Hastie, Manning, and Vaithyanathan, 2004; Dave et al., 2003; Matsumoto, Takamura, and Okumura, 2005; Pang and Lee, 2004; Pang et al., 2002; Turney, 2002).

This level of analysis assumes that each document contains opinions about a single object (a particular movie or product) from a single opinion holder. Thus, it is not applicable to documents which evaluate or compare multiple entities.

The task at the sentence level is to determine whether each sentence within a particular document (review) expresses a positive, negative, or neutral opinion (Khan and Baharudin, 2011a, 2011b; Meena and Prabhakar, 2007; Wilson, Wiebe, and Hoffmann, 2005). In other words, the sentence level classification considers each sentence as a separate unit and assumes that sentence should contain only one opinion. Moreover, such text analysis might still not be enough as a single sentence may contain different opinions about different facets of the same product.

Neither the document level nor the sentence level analyses discover what exactly people prefer and do not prefer. Instead of looking at language constructs (documents, paragraphs, sentences, clauses or phrases), the aspect level looks directly at the opinion itself. It is based on the idea that an opinion consists of a sentiment (positive or negative) and a target (of opinion). Aspect level was previously called *feature level* (Hu and Liu, 2004a). The goal of feature level classification is to produce a feature-based opinion summary of multiple reviews.

Thus, this research aims at making the classification at aspect level, in order to determine the most important features customers are interested in, and to use them in building a feature-based product search engine. In addition to develop a feature-based opinion summary of multiple reviews of each product page.

Since the sentiment classification problem is considered as a text classification problem, any existing supervised learning method can be applied, e.g., Naive Bayes (NB) classification (Moraes, Valiati, and GaviãO Neto, 2013; Pang et al., 2002), Maximum Entropy (ME) (Pang et al., 2002), and support vector machines (SVM) (Gamon, 2004; Moraes et al., 2013; Pang et al., 2002), etc. The effectiveness of machine learning techniques when applied to sentiment classification tasks is evaluated in the pioneering research by Pang et al. (2002). Theirs was the first paper to take this approach to classify movie reviews into two classes, positive and negative. The experimental results on the movie-review dataset produced via NB, ME, and SVM are substantially better than those results obtained through human generated baselines.

But their performance is not as remarkable as when they are used in topical text classification. Therefore, various kinds of feature sets, such as part-of-speech (POS) based features (Hatzivassiloglou and Wiebe, 2000), higher-order n-grams (Dave et al., 2003; Joshi and Penstein-Rosé, 2009; Pang et al., 2002), and word pairs and dependency relations (Dave et al., 2003; Gamon, 2004; Joshi and Penstein-Rosé, 2009; Subrahmanian and Reforgiato, 2008), and learning algorithms were tried by a large number of researchers to improve sentiment classification performance (Moraes et al., 2013).

However, the performance achieved by various methods is difficult to judge, since each method uses a variety of resources for training and different collections of documents for testing. Although the supervised techniques can achieve reasonable effectiveness, preparing training examples is time consuming. In addition, the effectiveness of the supervised techniques greatly depends on the representativeness of the training examples. In contrast, unsupervised approaches automatically extract product features from customer reviews without involving training examples. Therefore, unsupervised approaches seem to be more flexible than the supervised ones for environments in which various and frequently expanding products get discussed in customer reviews.

An opinion summary, also called an *aspect-based summary* or *feature-based summary*, was proposed in (Hu and Liu, 2004a). However, an opinion summary is quite different from a traditional single document or multi-document summary as an opinion summary is often centered on entities and their aspects and sentiments about them, and also has a quantitative side, while, the traditional single document summarization produces a short text from a long text by extracting some important sentences and discarding repeated information.

There are two ways to present the opinion summary: *structured summary*, or *short text summary*. The structured summary can also be visualized (Liu, Hu, and Cheng, 2005), where the user can clearly see the strengths and weaknesses of the product and its aspects in the minds of current consumers. In this research we will use the visual summarization.

2.2. Mining and Summarizing Chinese Opinions

Cross-language opinion analysis indicates to perform opinion analysis in multiple languages (Liu, 2006). Most of the research to date has focused on English and there are no effective resources or tools in other languages that can be used to build good sentiment classifiers in these languages. Therefore, the natural question is whether it is possible to use the automated machine translation capability and existing sentiment analysis resources and tools available in English to help build sentiment analysis systems in other languages.

Generally speaking, there are two unsupervised scenarios for borrowing English resources for sentiment analysis in other languages: one is to generate resources in a new language by relying on the resources available in English via cross-lingual projections, and then perform sentiment analysis in the English language based on the generated resources, which has been investigated by Mihalcea, Banea, and Wiebe, (2007); the other one is to translate the texts in a new language into English texts, and then perform sentiment analysis in the English language, which has been investigated by Wan (2008).

Mihalcea et al.'s study (2007) explores cross-lingual projections to generate subjectivity analysis resources in Romanian by leveraging on the tools and resources available in English. They have investigated two approaches: a lexicon-based approach based on Romanian subjectivity lexicon translated from English lexicon, and a corpus-based approach based on a Romanian subjectivity-annotated corpora obtained via cross-lingual projections.

Wan (2008) research exploited English sentiment resources to improve Chinese reviews sentiment analysis. Rather than simply projecting English resources onto Chinese resources, his approach first translates Chinese reviews into English reviews using multiple translators; it then uses a lexicon-based approach to classify each translated English version. The proposed approach consists of the following sentiment lexicons: positive lexicon, negative lexicon, negation lexicon, and intensifier lexicon. The algorithm then sums up the sentiment scores of the terms in the review considering negations and intensifiers. If the final score is less than 0, the review is negative, otherwise it is positive.

Furthermore, his approach performs sentiment analysis for both Chinese reviews and English reviews, and then uses both methods to combine the individual analysis results.

In contrast, the study on English sentiment analysis started much earlier and had a number of related corpus and resources which could provide many seed words with correct sentimental categories, while the study on Chinese sentiment analysis research started very late and the available resources on sentiment lexicon are extremely rare (Su and Li, 2012). There are two major problems that face the construction of a sentiment lexicon in Chinese: (1) Chinese words are very ambiguous, which makes it hard to compute the sentiment orientation of a word; (2) given the related research on sentiment analysis, available resources for constructing Chinese sentiment lexicons remain weak.

Generally, the main methods for constructing a sentiment lexicon can be categorized into three types (Su and Li, 2012): (1) Use existing electronic dictionary or word knowledge database to generate a sentiment lexicon. In the study of English, the resource of WordNet is popularly used (Esuli and Sebastiani, 2005; Hu and Liu, 2004a), while in Chinese study, the resource of HowNet is often used (Li, Yu, and Chen, 2011; Zhu, Min, Zhou, Huang, and Li-de, 2006). The main idea of this kind of approach is to find sentiment words in the dictionary which have similar semantics with the unknown word and then infer the sentiment orientation of the unknown words. However, this method cannot cover the context of the words. (2) Apply an unsupervised machine learning method: the co-occurrence frequency is often employed in a corpus to infer the word's close connection with some kind of polarity categories. For these approaches, the initial seed words play a central role in the success of constructing a high-precision sentiment lexicon. (3) Use human-annotated corpus: this inferred the sentiment tendency of one word according to the co-occurrence relationships or semantic relations on the basis of annotated sentiment classification corpus. This kind of method needs a larger amount of manually annotated corpus.

Unlike all the above studies, Su and Li (2012) in their work proposed a method that combines the English seed words and a bilingual statistics resource to build a Chinese sentiment lexicon which could better cover contextual information in a particular domain.

Research work focusing on Chinese sentiment analysis includes (Li and Sun, 2007; Tsou, Yuen, Kwong, and Lai, 2005; Wang, Wei, Li, Zhang, and Li, 2007; Ye, Shi, and Li, 2006). Such work represents heuristic extensions of the unsupervised or supervised methods for English sentiment analysis.

2.3. Feature-Based Product Ranking

There have been relatively few studies focusing on product feature-based ranking using customer reviews. The most closely related works on product ranking based on customer reviews are as follows:

Ghose and Yang (2009) proposed a ranking mechanism that ranks the reviews according to their helpfulness and their impact on sales. Their experiment results showed that subjectivity analysis can give useful clues about the helpfulness of a review and about its impact on sales.

Zhang et al. (2010) proposed a feature-based product ranking mechanism, that first identifies product features within a product category and analyze their frequencies and relative usage, then identifies subjective and comparative sentences and assigns sentiment orientations to these sentences. A weighted and directed feature graph has then been built by using statistics of all review sentences. The method employed a keyword strategy to identify feature sentences and the evaluation is carried out by a pRank algorithm using Amazon.com data.

Yu, Zha, Wang, and Chua (2011) developed an aspect ranking algorithm to identify the important aspects by simultaneously considering the aspect frequency and the influence of consumers' opinions given to each aspect on their overall opinions. The proposed system aimed to identify important aspects of a product from consumer reviews automatically.

Kong et al. (2011) proposed a feature-based product ranking technique by grouping users into familiar users (friends) and unfamiliar users (strangers), and assigning different weights to them based on their reliability degree, where friends on the top of the list are expected to be more reliable than the rest. Then, feature-based ranking list is calculated taking users' weighting into consideration.

Zhang, Cheng, Liao, and Choudhary (2012) introduced a product ranking system based on consumer review, considering the reviews credibility and posting date. This model has been built to help consumers to draw a conclusion about the quality of a particular product by developing a filtering mechanism to remove sentences/ comments unrelated to the product itself from consumer reviews, such as supported services.

Ghose, Ipeirotis, and Li (2012) proposed a utility-based ranking mechanism on product search engines (hotels) that incorporates multidimensional consumer preferences and social media signals. Their work aimed to highlight the tight linkages between user behavior on social media and search engines, by illustrating how social media can be mined and incorporated into a demand estimation model in order to generate a new ranking system in product search engines. The data was collected over three months from Travelocity.com.

Saranya (2013) proposed a ranking system that mainly focuses on identifying the product features from customer reviews automatically instead of identifying them manually. The proposed system aims to enhance the existing ranking systems and increase the ranking accuracy.

Our work differs from other proposed works in the following three aspects: first, this research is targeting Chinese language. Second, our work considers both product quality features and the associated services features because the consumers are not only looking for product quality but also the associated services to satisfy their requirements. Third, our work allows the searching for a particular product feature within the customers reviews and produces a ranked list based on this particular product feature.

3. Methodology

The main steps of the proposed work are summarized in Fig.1. Due to the absence of a benchmark database for a product search engine, it becomes crucial for us to build our own database that can be used for the development of the proposed system "Tsearch". The first step was crawling the needed data from the web. The database we have created was crawled from the leading Chinese e-commerce website: www.taobao.com.

The crawled data for each product mainly contained the product title, images, ID, metadata (price, shipping cost,...etc.) and all customers reviews for this product.

After collecting the data, a product title-based search index was built in order to enable fast searching for the products. The products search result returned from the search engine are then scored based on the proposed ranking function.

To rank the product search result, first we need to score each product based on product features extracted from customer reviews. We considered here two types of features: product-dependent features such as the writing speed of a memory card storage, and independent features such as the speed of the delivery service. The reviewers' opinions about the extracted features could be positive or negative.

Also, a front-end and back-end website was built for easy site usage and administration. The main homepage contains an input field which the user can use to enter his/her search query then hits a "Search" button. The search query will be submitted to the server-side of the system where it can be analyzed.

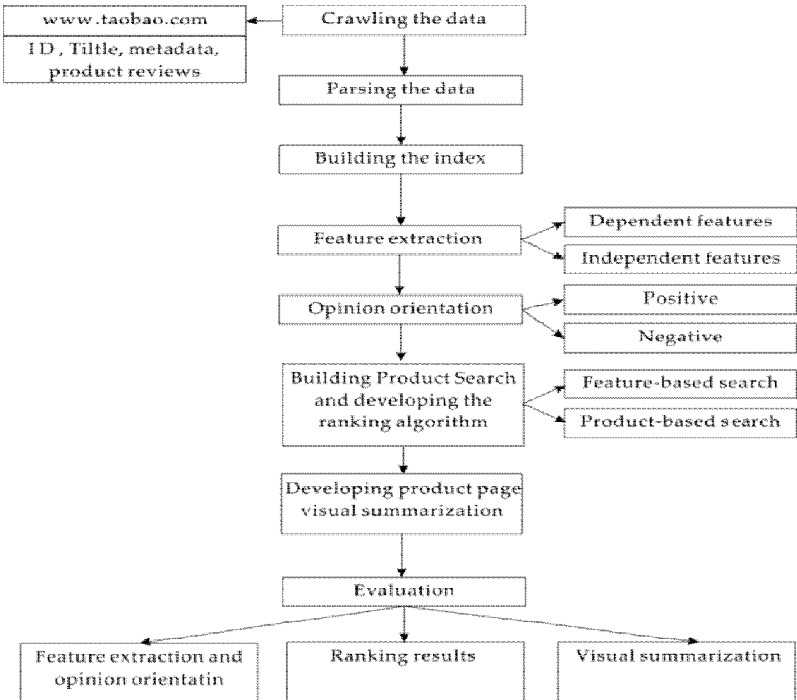


Figure 1: Methodology

The next step is generating the ranked products list of the matching products items and sending it back to the user. Each product page contains an aspect-based product features summarization of current customers opinions of the different product aspects.

The last step is the evaluation of the performance of the proposed system, and the effectiveness of feature extraction and opinion orientation, and the usability of the aspect-based visual summarization. In order to determine the effectiveness of feature extraction and sentiment classification we have used accuracy, precision and recall indexes on 2000 dependencies (feature/opinion pairs) selected randomly for 16 products from 5 domains. For the evaluation of the results obtained from the proposed ranking system and the aspect-based visual summarization, we have conducted a user survey. The participants are asked to assign a score of one (least satisfaction) to five (highest satisfaction), according to their satisfaction of the feature-based first page search result and also the visual summarization of each visited product page.

4. The Proposed Approach

4.1. System Architecture

An overview of the proposed system architecture is shown in Fig. 2. The system starts by crawling product pages from the web (www.taobao.com), parsing them to extract the product ID, title, images, metadata and all reviews, and then building a title-based search index. At the query time, the index is used for efficient retrieval of the matched products.

The ranking is generated according to the relevance score of each product which is calculated based on the product reviews. One of the main contributions of this work is allowing the search of products with the association of a particular product feature extracted from customers reviews, for example searching for Canon, with the mention of usability, image stabilization and color reproduction features. The proposed system supports a comma separated query where the first part of the query is the search keyword and the rest of it is the product features the user is interested in.

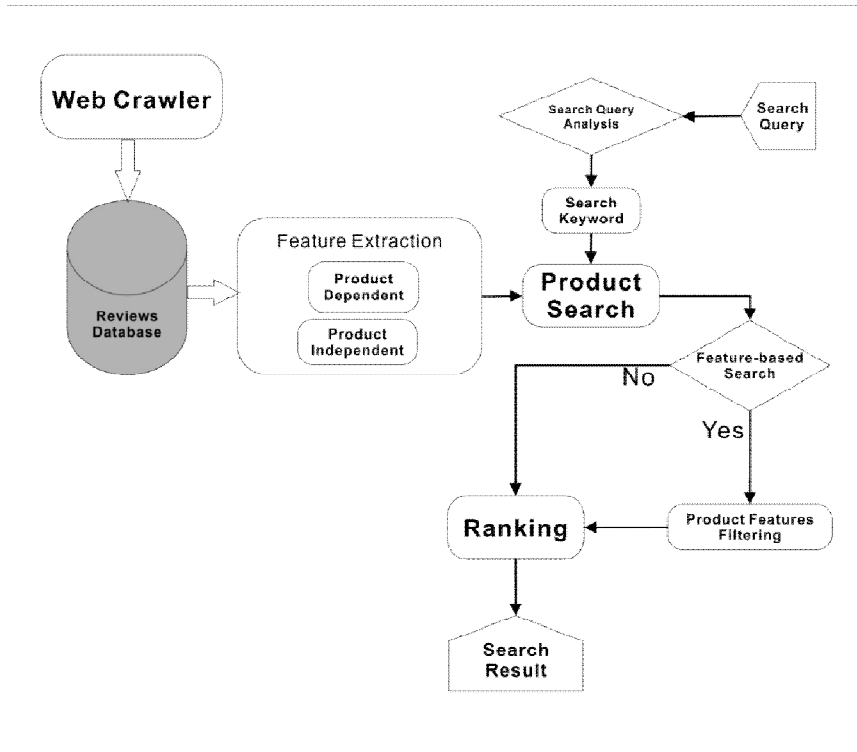


Figure 2: The Proposed System Architecture

To rank the product search result, first we need to score each product. The scoring is mainly based on the product features which refer to the aspects and specifications of the product that the reviewers have expressed their opinions on. These opinions could be positive or negative. In the review text, each opinion word with one product feature forms a *product feature opinion pair*. Various product feature extraction techniques have been proposed by many researchers (Hu and Liu, 2004a; Kobayashi, Iida, Inui, and Matsumoto, 2005; Popescu and Etzioni, 2005; Ravi Kumar and Raghuvver, 2013; Wang et al., 2013). The method proposed in (Hu and Liu, 2004a) has been seen as the elementary method for extracting product features. In our approach, the implemented method is similar to the one proposed in (Ravi Kumar and Raghuvver, 2013) which basically relies on extracting the typed dependencies and the collapsed dependencies for each opinion word in the review text. In our approach, Stanford typed and collapsed dependencies which have been applied in many fields were selected for product features extraction.

4.2. Stanford Typed Dependencies

The Stanford dependencies provide a representation of grammatical relations between words in a sentence that can be easily understood and effectively used by people without linguistic expertise who want to extract textual relations (De Marneffe & Manning, 2008). Stanford dependencies (SD) are triplets: the name of the relation, the governor and the dependent. Dependency relations (also known as binary relations) connect pairs of words or phrases, and name the relationship between these parts. For example: *nsubj (good, mobile)*, where *nsubj* (nominal subject) is the name of the dependency, and the ordering of the dependency is such that the head or the governor is listed first (in the example, the head is *good*), and the dependent is listed second (*mobile*).

Stanford dependencies have been successfully used by researchers in different domains: in textual entailment (Androutsopoulos & Malakasiotis, 2010), as well as in biomedical text mining (Kim, Ohta, Pyysalo, Kano, and Tsujii, 2009), information extraction (Wu and Weld, 2010); and sentiment analysis (Meena and Prabhakar, 2007; Zhuang, Jing, and Zhu, 2006). In addition to English, there is a Chinese version of Stanford dependencies (Chang, Tseng, Jurafsky, and Manning, 2009), which is also useful for many applications, such as Chinese sentiment analysis (Wu, Zhang, Huang, and Wu, 2009, 2011; Zhuang et al., 2006) and relation extraction (Huang, Sun, and Feng, 2008). Although there are several variants of Stanford dependencies for English, so far only a basic version (i.e. dependency tree structures) is available for Chinese. The wide variety of such Chinese-English differences includes the ordering of head nouns and relative clauses, and the ordering of prepositional phrases and the heads they modify.

Table 1: Chinese Grammatical Relations (Chang et al., 2009)

abbreviation	short description
nn	noun compound modifier
punct	punctuation
nsubj	nominal subject
conj	conjunct (links two conjuncts)
dobj	direct object
advmod	adverbial modifier
prep	prepositional modifier
nummod	number modifier
amod	adjectival modifier
pobj	prepositional object
rcomod	relative clause modifier
cpm	complementizer
assm	associative marker
assmod	associative modifier
cc	coordinating conjunction
clf	classifier modifier
ccomp	clausal complement
det	determiner
lobj	localizer object
range	dative object that is a quantifier phrase
asp	aspect marker
tmod	temporal modifier
plmod	localizer modifier of a preposition
attr	attributive
mmod	modal verb modifier
loc	localizer
top	topic
pccomp	clausal complement of a preposition
etc	etc modifier
lccomp	clausal complement of a localizer
ordmod	ordinal number modifier
xsubj	controlling subject
neg	negative modifier
rcomp	resultative complement
comod	coordinated verb compound modifier
vmod	verb modifier
prtmod	particles such 所, 以, 来, 而
ba	"ba" construction
dvpm	manner DE (地) modifier
dvpmod	a "XP+DE" (地) phrase that modifies VP
prnmod	parenthetical modifier
cop	copular
pass	passive marker
nsubjpass	nominal passive subject

Table 2: Penn POS

POS	POS description	Example
AD	adverb	还
AS	aspect marker	着
BA	把 in ba-construction,	把, 将
CC	coordinating conjunction	和
CD	cardinal number	一百
CS	subordinating conjunction	虽然
DEC	的 in a relative-clause	的
DEG	associative 的	的
DER	得 in V-de const. and V-de-R	得
DEV	地 before VP	地
DT	determiner	这
ETC	for words 等, 等等	等, 等等
FW	foreign words	ISO
IJ	interjection	啊
JJ	other noun-modier,	男, 共同
LB	被 in long bei-const,	被, 给
LC	localizer	里
M	measure word	个
MSP	other particle	所
NN	common noun	书
NR	proper noun	美国
NT	temporal noun	今天
OD	ordinal number	第一
ON	onomatopoeia,	哈哈, 哗哗
P	preposition excl. 被 and 把	从
PN	pronoun	他
PU	punctuation	'?。'
SB	被 in short bei-const,	被, 给
SP	sentence-nal particle	吗
VA	predicative adjective	红
VC	是	是
VE	as the main verb	有
VV	other verb	走

The Chinese version of Stanford dependencies contains 45 different dependencies, 44 are grammatical relations, and a default 45th relation dep (dependent). If a dependency matches no patterns, it will have the most generic relation dep. The description of the 44 grammatical relations are listed in Table 1.

The Chinese typed dependencies make use of Penn POS tags (Chang et al., 2009). As shown in Table 2, Penn POS tag set has 33 tags distributed as follows:

Verb, adjective(4)	VA, VC, VE, VV	Noun (3)	NR, NT, NN
Localizer (1)	LC	Measure word (1)	M
Adverb (1)	AD	Pronoun (1)	PN
Determiner and number (3)	DT, CD, OD	Preposition (1)	P
Conjunction (2)	CC, CS	Others (8)	IJ, ON, PU, JJ, FW, LB, SB, BA
Particle (8)	DEC, DEG, DER, DEV, SP, AS, ETC, SP, MSP		

4.3. Product Features Extraction

The classification of the product features may differ from one approach to another. In our approach we classify the product features into two categories: product-dependent features and product-independent features. The product-dependent features are the features that describe the product itself or its components such as the battery life of smart cellphone and the writing speed of a USB flash disk. The product-independent features are the features that describe the associated services, such as the quality of the packaging material, speed of delivery service and after sale support. In our approach, we assign different weights to each category, assuming that the customers are likely to be more concerned about the product-dependent features more than product-independent features.

The first step in the processing of customer reviews is sentence segmentation. In Chinese language, there is no capitalization or space between words, which makes a Latin-based sentence segmenter produce unexpected results. Therefore, we developed a punctuation-based sentence segmenter for this task. After segmenting each review text into individual sentences, we start parsing and tagging each word in the sentence as shown in the Fig.3. We used Stanford parser (Chang et al., 2009; Levy and Manning, 2003) (which support Chinese language) for part of speech (POS) tagging and dependencies extraction. The dependencies produced by the parser were used as an initial estimation for extracting the product features opinion pairs. The first argument in each dependency is the opinion word, and a list of positive and negative opinion words are used as seeds for sentiment classification.

```

Your query
上网#过了是正品。
Segmentation
上网 查过了 是 正品 。
Tagging
上网/VV 查过/VV 了/AS 是/VC 正品/NN 。 /PU
Parse
(ROOT
 (IP
 (IP
 (VP (VV上网)
 (VP (VV查过) (AS了))))
 (VP (VC是)
 (NP (NN正品)))
 (PU。 )))
Typed dependencies
mmod(查过-2, 上网 1)
conj(是-4, 查过-2)
asp(查过-2, 了 3)
root(ROOT-0, 是-4)
attr(是-4, 正品-5)
Typed dependencies, collapsed
mmod(查过-2, 上网 1)
conj(是-4, 查过-2)
asp(查过-2, 了 3)
root(ROOT-0, 是-4)
attr(是-4, 正品-5)
top(是-4, 查过-2)

```

Figure 3: Example of a review for an iPhone, (the translation into English is: checked on the Internet it is original), the dependency attr contains the positive product features "it is original"

The product features extraction algorithm is listed in Algorithm.1. This algorithm aims at extracting the candidate feature opinion pairs based on using different combinations of dependencies. Then based on POS tags of the first and second arguments, we select the appropriate pairs to be candidates for scoring. Five different dependencies are selected in our work as follows:

Algorithm Inputs:

S: review sentence.

D(*gov*, *dep*): a dependency in the sentence **S** consist of two parts *gov* and *dep*

foreach *Dependency D_i* **in** **S** **do**

if *D_i* == *nsubj* **then**

if *POS*(*gov*) == *NN* & *POS*(*dep*) == *CD* **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *POS*(*gov*) == *VV* & *POS*(*dep*) ∈ {*NN*, *VA*, *NR*} **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *POS*(*gov*) == *VA* & *POS*(*dep*) == *NN* **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *D_i* == *doobj* **then**

if *POS*(*gov*) == *VE* & *POS*(*dep*) == *NN* **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *D_i* == *ccomp* **then**

if *POS*(*gov*) == *VC* & *POS*(*dep*) == *VA* **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *D_i* == *nn* **then**

if *POS*(*gov*) == *NN* & *POS*(*dep*) == *NN* **then**

 | **ScoreFeature**(*TEXT*(*gov*))

else if *D_i* == *attr* **then**

if *TEXT*(*gov*) == 是 & *POS*(*dep*) == *NN* **then**

 | **ScoreFeature**(*TEXT*(*dep*))

end

end

Algorithm 1: Candidate Product Features Selection Algorithm

- 1) **nsubj - Nominal Subject**: A nominal subject is a noun phrase which is the syntactic subject of a clause. The governor of this relation might not always be a verb: when the verb is a copula verb, the root of the clause is the complement of the copula verb, which can be an adjective or noun (De Marneffe and Manning, 2008). The Chinese copula verb (the equivalent of English "to be" and all its forms) is 是 "shì".

However, it is normally only used when its complement is a noun or noun phrase. Prepositions and predicate adjectives function as verbs themselves, so in sentences where the predicate is an adjectival or prepositional phrase, 是“shì” is not required. Five nominal subject combinations are used for product features candidate selection as follows:

- a) nsubj(NN,CD): for example (很多, 礼品)
 - b) nsubj(VV,NN): for example(全真, 配件)
 - c) nsubj(VV,VA): for example(不错, 手机)
 - d) nsubj(VV,NR): for example(喜欢, 外观)
 - e) nsubj(VA,NN): for example(稳定, 信号)
- 2) **dobj- Direct Object:** The direct object of a verb phrase (VP) is the noun phrase which is the (accusative) object of the verb (De Marneffe and Manning, 2008). A direct object will follow a transitive verb (a type of action verb). Direct objects can be nouns, pronouns, phrases, or clauses. Only action verbs can have direct objects. If the verb is linking, then the word that answers the what? or who? question is a subject complement. One direct object dependency combination is used for product features candidate selection as follows:
- a) dobj(VE,NN): for example(没有, 瑕疵)
- 3) **ccomp- Clausal Complement:** There are more ccomp in the Chinese sentences and less in English. A clausal complement of a verb or adjective is a dependent clause with an internal subject which functions like an object of the verb, or adjective. Clausal complements for nouns are limited to complement clauses with a subset of nouns like “fact” or “report”. Such clausal complements are usually finite (though there are occasional remnant English subjunctives) (De Marneffe and Manning, 2008). One clausal complement dependency is used for product features candidate selection as follows:
- a) ccomp (VC,VA): for example (是, 新的)
- 4) **nn-Nominal Compound Modifier:** A noun compound modifier of a noun phrase (NP) is any noun that serves to modify the head noun (De Marneffe and Manning, 2008). Suppose the product feature is of more than one word then we can find this by using this dependency. One nominal compound modifier dependency is used for product features candidate selection as follows:
- a) nn(NN,NN): for example (声音大, 震动)
(模糊, 照相)

(非常好, 质量)

- 5) **attr- Attributive**: In grammar, an attributive is a word or phrase within a noun phrase that modifies the head noun. It may be an: attributive adjective, attributive noun, attributive verb, or other part of speech, such as an attributive numeral. For product features candidate selection, one attributive dependency relation is selected as follows:

- a) attr(VC,NN): for example (是, 正品)
(是, 新的)

As shown in Fig.4, the selected typed dependencies happened to have high frequencies compared to other typed dependencies. This explains why we only select these five typed dependencies. The result shown in Fig.4 was obtained by collecting more than **537,638 reviews** which include more than **21.9 million** typed dependency tuples.




Figure 4: Dependency Frequency Distribution, the Blue Bars Indicate the Typed Dependencies

4.4. Product Features Scoring

The candidate features extracted using the previously mentioned algorithm listed in Algorithm.1 may not always truly represent product feature pairs. To ensure that, each candidate product feature must contain an opinion word to be selected as a product feature pair. In the *nsubj*, *dobj*, *ccomp* and *nn* typed dependencies, the dependent (the first argument) is the opinion word, while in the *attr* typed dependency, the governor (the second argument) is the opinion word.

The scoring algorithm is listed in Algorithm.2. According to this algorithm, the scoring is calculated by searching for the opinion word in the predefined seed lists: a positive opinion list (POL), and a negative opinion list (NOL). These seed lists are constructed using a Chinese sentiment dictionary, the HowNet- VSA (Dong and Dong, 2006). It contains 3745 negative opinion words and 4534 positive opinion words. If the opinion word is in the POL, we select the candidate feature pair and give a score of 1, and if the opinion word is in the NOL, we select the candidate feature pair and give a score of -1, otherwise we give a score of 0 if the opinion word does not exist in POL and NOL, which is equivalent to the rejection of the candidate product feature pair.

4.5. Product Ranking

When consumers enter the search query into a product search engine, it returns a list of the matched products in a random order. To rank this list, many numerical factors could be used to score each product such as the price of the products. In our work, we propose a new scoring function which is based on the customer reviews.

First we consider the products which have nonzero value for all product features pairs should have a high ranking score regardless of the value of the feature product pairs. We refer to this as the product features completeness *PC*. So, for a given product category *i* (e.g. smartphone), if the total number of product features pairs for *i* equal to N_i , then the *PC* for the product page *j* given the category *i*, PC_{ij} will be:

$$PC_{ij} = \frac{\text{Count}(P_{ij}, F)}{N_i} \quad (1)$$

where:

count(P_{ij} , F) is the total number of product feature pairs F which the product page P_j have given the product category i .

Algorithm Inputs:

$D(gov, dep)$: a tuple consists of dependency name D , governor gov and dependent dep

POL: Positive Opinion List.

NOL: Negative Opinion List.

Algorithm Output:

Score: The score of the input dependency.

if $D \subset \{nsubj, dobj, ccomp, nn\}$ **then**

 | $OpinionWord = \text{TEXT}(gov)$

else

 | $OpinionWord = \text{TEXT}(dep)$

end

if $OpinionWord \in POL$ **then**

 | **Score** = 1

else if $OpinionWord \in NOL$ **then**

 | **Score** = -1

else

 | **Score** = 0

end

Algorithm 2: Product Features Pairs Scoring Algorithm used in Our System

The second factor in the scoring function is the product page score PS_{ij} which we calculate as follows:

$$PS_{ij} = \frac{1}{M} \sum_{\tau=1}^M \frac{C(P_j, F_{\tau})}{\max(C(P_i, F_{\tau}))} \quad (2)$$

while:

$$C(P_j, F_{\tau}) = \begin{cases} \text{Count}(P_j, F_{\tau}^{pos}) - \text{Count}(P_j, F_{\tau}^{neg}) & \text{if } \text{Count}(P_j, F_{\tau}^{pos}) > \text{Count}(P_j, F_{\tau}^{neg}) \\ 0 & \text{if } \text{Count}(P_j, F_{\tau}^{pos}) \leq \text{Count}(P_j, F_{\tau}^{neg}) \end{cases} \quad (3)$$

where:

$C(P_j, F_\tau)$ is the value of the product feature pair τ in the product page j ; $\max(C(P_i, F_\tau))$ is the maximum value of the product feature pair τ in the all product pages P in the category i ; M is the total number of product features pairs within all products pages; $Count(P_j, F_{pos}^\tau)$ is the total number of positive product feature pairs F_{pos}^τ ; $Count(P_j, F_{neg}^\tau)$ is the total number of negative product feature pairs F_{neg}^τ .

When the customer provides a search query without specifying any product feature, then we calculate the rank of the product j ; P_jR according to the following function:

$$P_jR_{query} = 0.3P_jC + 0.5P_jS_{dependent} + 0.2P_jS_{independent} \quad (4)$$

If the customer provides the product feature he/she is interested in, then the product score P_jR is calculated using only these product features:

$$P_jR_{query} = 0.3P_jC + 0.7P_jS_{user} \quad (5)$$

5. System Implementation

The implementation of the proposed system is based on system architecture shown in Fig.2.

5.1. Data Collection

The data was crawled from the leading e-commerce website in China *taobao.com*. The details of the data set are given in Table 3. This data set contains consumer reviews on 16 popular products from 5 domains (phone, digital camera, rice cooker, soymilk maker and laptop). Each product belongs to a single product category. Each product category contains a set of the top three popular products (Brands) preferred by taobao.com consumers within that category.

However, more than 16,071 product pages with more than 537,638 reviews have been crawled, and the following information has been extracted:

- 1) Product title: which we used to build the search index.
- 2) Product ID: each product has a unique ID where all relevant information can be crawled based on this ID. The product ID can also be used for future updates of the product reviews.
- 3) Product metadata: such as the price, delivery cost, the number of deals within the last 30 days and the number of reviews.
- 4) Product reviews: all product reviews with their supplementary information such as: posting date and review usefulness.

Table 3: Statistics of the Data Sets

Domain	Brand	Consumer	Product name	Product pages	Reviews	Sentences	Max reviews number
phone	MIUI	16%	M4	2779	73887	334375	33620
	Nokia	11%	Lumia 1020	361	15332	69338	3168
	Apple	8%	iPhone 6	4568	41866	191721	9956
	Samsung	8%	Galaxy Note 3	3885	97653	440224	1567
digital camera	Casio	33%	EX-ZR1200	527	25349	117571	9267
	Canon	29%	SX510 HS	298	16419	82082	8493
	Sony	10%	DSC-HX300	353	5033	23629	1051
rice cooker	Midea	23%	FS406C	778	52758	245194	13729
	Supor	14%	CFXB40FC118-75	282	37813	176649	28107
	Joyoung	7%	JYF-40FS06	119	15716	71263	1051
soymilk maker	Joyoung	67%	DJ13B-D08D	387	94642	432480	48825
	Midea	16%	DE12G13	536	30540	139387	16218
	Peskoe	7%	D09	431	18733	84494	2313
laptop	Lenovo	23%	G510AM-IFI	352	6066	29096	2936
	Asus	17%	X550	189	3182	15340	1575
	Dell	12%	5R(5537) INS15R-5528	226	2649	12171	664

5.2. Data Indexing

After collecting the data, we have built a title-based index for all crawled products. This index serves as the second step in building our product search engine after crawling. It is based on the product title because the product title contains all basic information related to the product, such as the product name, brand and model. Also, the users use a keyword which can always be found in the product title. The search index was designed to be expandable, which allows any updates for existing products or new products to be easily added.

The search index is always implemented as a word-based index, which means that each word is a key in the index. When the user query contains more than one word, a logical AND is used between the search result to combine the duplicated results. But, while Chinese language is a character-based language and there is no separation between the characters, it is easier if we index product titles using only individual characters as an indexing key.

In addition to the Chinese characters, the product title always contains English words, such as the brand name, or contains numbers, such as the model number, which makes the building of the search index more complicated. So, before building the search index, each title should be segmented to separate Chinese characters from English words and numbers. The segmentation can be easily done if we consider the coding set for each language, where we can convert the title text to unicode format, then based on the character code, we can tell whether it is Chinese or English. Any sequence of English characters are grouped to form a word, and the space and other punctuations are used to separate English words.

5.3. Search Query Analysis

We have developed our search engine to be able to accept a search keyword written in Chinese, English or both. But, in addition to the search keyword, the users can also input an unlimited number of product features he/she is interested in. The product features can be entered by adding a comma “,” after the search keyword. If the user is looking for more than one product feature, he/she can also use a comma “,” to separate different product features.

5.4. Search Result Listing

When the user submits the search query, the system responds with all products that contain the required search keywords. As shown in Fig.5, the search result is presented as a grid, where each cell contains the product image, title and the ranking score. Also, the search keywords are highlighted within the product title. If the search result contains more than 30 products, a pagination control automatically appears at the bottom of the page with a navigation buttons to allow the user to browse all the search results.



Figure 5: Search Result as a Grid Layout Based on the Scores Produced by our Ranking Function

5.5. View Product Features Summarization

When the user clicks on the title or the image of the product in the search results page, the product details page will be opened. The product metadata with the image and the title will appear at the top of the page. In addition to that, a visual summary of (a maximum of) the top 10 product features will be displayed in the right side of the product image as a bar graph, as shown in Fig.6. This summary is very useful for users who just want to make sure that a specific product feature has received a high number of positive reviews or to have a quick look at the opinions about the other product features without going into details.

If the user wants to have more information about the product features mentioned in the reviews, he/she can sort all the reviews, where the review contains more features



Figure 6: The product details view which in addition to the product image, title and metadata contains a summary bar graph of the top 10 product features

The screenshot shows a web browser window with a search bar and a review text. The review text is: "卖家服务好, 手机原装没激活, 仔细检查完美无瑕, 物流也快, 原以为赶上11.11会耽搁几天, 没想到发货第二天就送到, 今天抽空自取, 验货正品!". Below the review text, a table displays the extracted dependencies. The table has four columns: Polarity, Feature, Dependency Type, and Score. The rows are grouped by feature, with the feature name highlighted in yellow. The features are: 卖家服务好, 手机原装没激活, 仔细检查完美无瑕, and 物流也快. Each feature group contains rows for the feature itself and its dependencies, with the feature row highlighted in green.

Polarity	Feature	Dependency Type	Score
卖家服务好			
服务/NN	卖家/NN	nn	0
好/VA	服务/NN	nsubj	1
ROOT/null	好/VA	root	0
手机原装没激活			
原装/NN	手机/NN	nn	1
激活/VV	原装/NN	nsubj	0
激活/VV	没/AD	neg	0
ROOT/null	激活/VV	root	0
仔细检查完美无瑕			
检查/VV	仔细/AD	advmod	0
ROOT/null	检查/VV	root	0
无瑕/NN	完美/NN	nn	1
检查/VV	无瑕/NN	doobj	0
物流也快			
快/VA	物流/NN	nsubj	1
快/VA	也/AD	advmod	0
ROOT/null	快/VA	root	0

Figure 7: The review details can be obtained by clicking on the review text which will show up a table contains all typed dependencies extracted from the review text with product features highlighted

or a particular feature that the user is interested in, it appears first, which saves the user time and effort.

The more detailed information about the product review and all product features within it can be viewed as a typed dependency by clicking on the review text itself. This will produce a table containing all types dependencies extracted from the review, with highlighted rows to indicate where the product feature is mentioned and whether it is positive or negative, as shown in Fig.7.

6. Evaluation

To evaluate our system, the evaluation process should be conducted in two levels: the first one to measure the accuracy of product feature extraction and classification, and the second one to determine the efficiency of the search results and the usability of visual summarization.

6.1. Evaluation of Feature Extraction

To evaluate the performance of our approach and to measure the quality of the retrieved product features, we have adopted accuracy, precision and recall indexes as shown in Table 4. Accuracy is the portion of all true predicted instances against all predicted instances. An accuracy of 100% means that the predicted instances are exactly the same as the actual instances. Precision is the portion of true positive predicted instances against all positive predicted instances. The higher the measure's value, the better the retrieval is.

Table 4: Confusion Matrix

	Predicted positives	Predicted negatives
Actual positive	TP	FN
Actual negative	FP	TN

Where, TP is the number of true positive instances, FN is the number of false negative instances, FP is the number of false positive instances, and finally TN is the number of true negative instances.

In order to implement the evaluation process, we have used 1000 dependencies predicted positive and 1000 dependencies predicted negative (both have been selected randomly) for each product. The results are shown in Table 5.

Table 5: Product Feature Evaluation

Product name	Domain	TP	FP	FN	TN	Accuracy	Precision	Recall
MIUI - M4	Phone	900	100	115	885	89.25	90	88.67
Nokia - Lumia 1020		890	110	108	892	89.1	89	89.18
Apple - iPhone 6		888	112	127	873	88.05	88.8	87.49
Samsung - Galaxy Note 3		890	110	111	889	88.95	89	88.91
Average of the phone domain						88.84	89.20	88.56
Casio - EX-ZR1200	digital	870	130	177	823	84.65	87	83.09
Canon - SX510 HS	camera	877	133	168	832	85.02	86.83	83.92
Sony - DSC-HX300		901	99	150	850	87.55	90.1	85.73
Average of the digital camera domain						85.74	87.98	84.25
Midea - FS406C	rice cooker	920	80	114	886	90.3	92	88.97
Supor - CFXB40FC118-75		917	83	114	886	90.15	91.7	88.94
Joyoung - JYF-40FS06		900	100	120	880	89	90	88.24
Average of the rice cooker domain						89.82	91.23	88.72
Joyoung - DJ13B-D08D	soymilk maker	923	77	111	889	90.6	92.3	89.26
Midea - DE12G13		913	87	120	880	89.65	91.3	88.38
Peskoe - D09		924	76	114	886	90.5	92.4	89.02
Average of the soymilk maker domain						90.25	92.00	88.89
Lenovo - G510AM-IFI	laptop	871	129	154	846	85.85	87.1	84.98
Asus - X550		865	135	163	837	85.1	86.5	84.14
Dell - 5R(5537) INS15R-5528		830	170	177	823	82.65	83	82.42
Average of the laptop domain						84.53	85.53	83.85
Total Average						87.84	89.19	86.85

On the basis of these evaluation measures, results showed that the proposed approach performed well in most of the domains with 87.84% accuracy, 89.19% precision, and 86.85% in the recall index, best performance was in soymilk maker domain with 90.25% accuracy, 92 % precision, and 88.89% in the recall index as shown in Table 5.

6.2. Evaluation of the Ranking Results

The use of customer reviews for product ranking is still a subjective problem, there has not been a commonly recognized method for validating a ranking system (Zhang, Cheng, and Liao, 2012). However, in our evaluation, we have used a user study for evaluating the feature based results.

However, in order to evaluate the feature-based ranking result, we conducted a user survey that examines consumers' responses to different personalized ranking mechanisms according to the features preferred by them. First of all, we randomly select one product name from each product category than we randomly select 20 product pages from each product name that have at least 20 reviews. We ask 40 volunteers to participate in our evaluation. Each participant was asked to enter a query that contains the name of the product followed by the product feature he/she is interested in. Then, they were asked to assign a score of one (least satisfaction) to five (highest satisfaction), according to their satisfaction of the proposed product search engine's first page results. The results are shown in Table 6. The satisfaction average of the ranking results is 3.92, which reflect the user satisfaction of the feature-based ranking approach.

Table 6: The Average Satisfaction Scores of the Customized Result

User Number	Score	User Number	Score	User Number	Score	User Number	Score
1	3.8	11	3.9	21	4.2	31	3.7
2	4	12	3.8	22	3.7	32	3.9
3	3.9	13	4.2	23	4.2	33	4
4	3.7	14	3.7	24	4	34	3.8
5	3.7	15	3.7	25	3.8	35	4.2
6	4.2	16	3.9	26	3.9	36	3.8
7	4.2	17	4	27	4	37	3.9
8	4	18	3.8	28	3.7	38	3.8
9	4.2	19	3.8	29	4.2	39	3.7
10	3.8	20	4.2	30	3.7	40	4.2
Average satisfaction 3.92							

6.3. Evaluation of Summarization

Only 19 participants from the above 40 participants are volunteered to make the evaluation of the usefulness of the visual summarization. Each participant randomly selects four products from the 20 product pages for each product category and read their customer reviews. Then, they were asked to assign a score of one (least satisfaction) to five (highest satisfaction), according to their satisfaction of the visual summarization of each visited product page.

Table 7 presents the average satisfaction scores for the visual summarization from each participant. The satisfaction average is 4.23, which reflects the user satisfaction of the proposed approach. In summary, we can see that the proposed technique is very promising, especially for sentence orientation prediction, and feature based ranking.

Table 7: The Average Satisfaction Scores of the Usefulness of Summarization

Category	phone				Digital Camera				Rice Cooker				Soymilk maker				Laptop				
Users pp	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	
User 1	4.5	4.3	4.5	4.2	3.9	4	3.9	3.9	4.6	4.5	4.6	4.5	4.7	4.7	4.7	4.6	4	3.9	3.8	4.1	
User 2	3.8	3.8	3.9	3.8	3.8	3.8	4	3.7	4.4	4.5	4.6	4.5	4.7	4.5	4.6	4.6	4.1	3.9	4	3.9	
User 3	4.5	4.5	4.2	4.4	3.7	4.1	4.1	4	4.5	4.5	4.4	4.4	4.5	4.6	4.6	4.5	4	3.8	3.9	4	
User 4	4	4	4	4	3.7	3.8	3.9	4	4.5	4.5	4.5	4.4	4.6	4.6	4.6	4.5	4.1	3.8	3.9	4.1	
User 5	3.9	4.1	4	4.1	4	4	4	4	4.3	4.4	4.4	4.6	4.5	4.4	4.5	4.7	3.9	4	4	3.9	
User 6	4.6	4.1	4.3	4.5	3.9	3.9	4	3.8	4.4	4.5	4.3	4.4	4.4	4.5	4.6	4.4	4	3.9	3.8	4.1	
User 7	4.3	4.4	4.2	4.2	3.8	4.1	4.1	4	4.3	4.4	4.5	4.5	4.6	4.4	4.5	4.6	4.1	3.8	4	4	
User 8	4.1	4.3	4.1	4.3	3.9	3.8	3.9	4	4.5	4.6	4.6	4.4	4.7	4.6	4.7	4.5	3.9	4	4	3.9	
User 9	4.2	4.5	4.3	4.4	3.9	4	4	3.8	4.5	4.6	4.6	4.6	4.7	4.6	4.7	4.6	4	3.9	3.9	3.8	
User 10	4.4	4.3	4.4	4.4	3.8	3.9	4.1	4	4.4	4.5	4.5	4.5	4.6	4.5	4.6	4.6	4.1	3.8	3.7	4	
User 11	4.3	4.1	4.4	4.2	3.8	4.1	4	3.9	4.6	4.5	4.3	4.4	4.3	4.7	4.6	4.5	3.9	3.7	4	3.9	
User 12	4.5	4.3	4.5	4.4	3.8	3.9	4.1	3.9	4.6	4.4	4.5	4.4	4.6	4.7	4.5	4.5	4.1	3.7	4	4.1	
User 13	4.4	4	4.3	4.2	4	4.1	4	3.9	4.5	4.6	4.5	4.6	4.6	4.6	4.7	4.7	4	3.8	3.7	3.8	
User 14	4.5	4.4	4.6	4.5	3.9	4	4	4	4.3	4.5	4.3	4.5	4.4	4.4	4.6	4.6	4	4	3.8	4	
User 15	4.3	4.4	4.2	4.1	3.8	3.9	3.9	3.8	4.4	4.4	4.4	4.4	4.5	4.5	4.5	4.5	4.1	3.9	4	3.9	
User 16	4.2	4.3	4.4	4.4	4	4.1	4	3.7	4.4	4.4	4.5	4.5	4.6	4.5	4.5	4.6	3.9	3.8	4	4.1	
User 17	3.8	3.9	3.9	4	3.9	3.9	4.1	3.8	4.5	4.5	4.3	4.4	4.4	4.6	4.6	4.5	4	4	3.9	4	
User 18	4.1	4.5	4.4	4.3	3.8	4	3.9	4	4.5	4.4	4.4	4.4	4.5	4.6	4.5	4.5	4	3.9	4	3.8	
User 19	4.3	4.4	4.3	3.3	4	4.1	4.1	4	4.6	4.6	4.5	4.5	4.6	4.7	4.7	4.6	3.9	3.8	3.8	4.1	
Average	4.24				3.94				4.47				4.57				3.94				
Total average																					4.23

7. Conclusion and Future Perspectives

In this paper, we proposed a ranking system for a product search engine. The data have been crawled from taobao.com. The ranking results considered all dependent and independent features in customers reviews in response to a customer queries. The evaluation of the proposed system shows a high level of accuracy in product feature extraction, and the satisfaction average of both the ranking system and the visual summarization from the participants range from 3.92 for the ranking results and 4.23 for the visual summarization, which reflects the user satisfaction and the effectiveness of the proposed approach.

Building an effective ranking system that satisfies the customers requirements is a challenging task.

This research was only focused on building a product search engine based on product features customers are interested in, while current customers generally based their purchasing decisions on different criteria, such as price, quality, speed of the delivery services, in addition to the availability of the product features' description the consumer is interested in within other customers reviews. Thus, future research into a ranking mechanism of product search engines that incorporates multidimensional consumer preferences and social media signals could lead to significant surplus gain for consumers, and could enhance the customer e-shopping experience.

This research adopted the five different Stanford typed dependencies for Chinese language: *nsubj*, *attr*, *nn*, *ccomp*, and *dobj* in order to determine the candidate product feature opinion pairs. Another area for future research is to examine the other Stanford typed dependencies in order to determine the presence probability of other candidates product feature opinion pairs within Chinese customers reviews.

The consideration of multiple typed dependencies to extract one product feature presents another area for future research.

References

- Abbasi, A., Chen, H., and Salem, A. (2008). Sentiment analysis in multiple languages: Feature selection for opinion classification in web forums. *ACM Trans. Inf. Syst.*, 26 (3), 12:1–12:34.
- Abulaish, M., Jahiruddin, Doja, M. N., and Ahmad, T. (2009). Feature and opinion mining for customer review summarization. In *Proceedings of the 3rd international conference on pattern recognition and machine intelligence* (pp. 219–224). Springer-Verlag.
- Aisopos, F., Papadakis, G., Tserpes, K., and Varvarigou, T. (2012). Textual and contextual patterns for sentiment analysis over microblogs. In *Proceedings of the 21st international conference companion on world wide web* (pp. 453–454). New York, USA: ACM.
- Aisopos, F., Papadakis, G., and Varvarigou, T. (2011). Sentiment analysis of social media content using n-gram graphs. In *Proceedings of the 3rd acm sigmm international workshop on social media* (pp. 9–14). New York, USA: ACM.
- Androutopoulos, I., and Malakasiotis, P. (2010). A survey of paraphrasing and textual entailment methods. *J. Artif. Int. Res.*, 38 (1), 135–187.
- Bafina, K., and Toshniwal, D. (2013). Features based summarization of consumers' reviews of online products. *Procedia Computer Science*, 22 , 142–151.

- Becker, I., and Aharonson, V. (2010). Last but definitely not least: On the role of the last sentence in automatic polarity-classification. In Proceedings of the acl 2010 conference short papers (pp. 331–335). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Beineke, P., Hastie, T., Manning, C., and Vaithyanathan, S. (2004). Exploring sentiment summarization. In Proceedings of the AAAI spring symposium on exploring attitude and affect in text: Theories and applications. AAAI Press.
- Chang, P.-C., Tseng, H., Jurafsky, D., and Manning, C. D. (2009). Discriminative reordering with chinese grammatical relations features. In Proceedings of the third workshop on syntax and structure in statistical translation (pp. 51–59).
- Das, S., and Chen, M. (2001). Yahoo! for amazon: Extracting market sentiment from stock message boards. In Proceedings of Asia pacific finance association annual conference (pp. 22–25).
- Dave, K., Lawrence, S., and Pennock, D. M. (2003). Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In Proceedings of the 12th international conference on world wide web (pp. 519– 528). New York, USA: ACM.
- Davidov, D., Tsur, O., and Rappoport, A. (2010). Enhanced sentiment learning using twitter hashtags and smileys. In Proceedings of the 23rd international conference on computational linguistics: Posters (pp. 241–249). Stroudsburg, PA, USA: Association for Computational Linguistics.
- De Marneffe, M.-C., and Manning, C. D. (2008). The Stanford typed dependencies representation. In Coling 2008: Proceedings of the workshop on cross- framework and cross-domain parser evaluation (pp. 1–8). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Dong, Z., and Dong, Q. (2006). hownet and the computation of meaning. River Edge, NJ, USA: World scientific Publishing Co., Inc.
- Esuli, A., and Sebastiani, F. (2005). Determining the semantic orientation of terms through gloss classification. In Proceedings of the 14th acm international conference on information and knowledge management (pp. 617–624). New York, USA: ACM.
- Gamon, M. (2004). Sentiment classification on customer feedback data: Noisy data, large feature vectors, and the role of linguistic analysis. In Proceedings of the 20th international conference on computational linguistics. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Ghiassi, M., Skinner, J., and Zimbra, D. (2013). Twitter brand sentiment analysis: A hybrid system using n-gram analysis and dynamic artificial neural network. *Expert Syst. Appl.*, 40 (16), 6266–6282.
- Ghose, A., Ipeirotis, P. G., and Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowd-sourced content. *Marketing Science*, 31 (3).
- Ghose, A., Ipeirotis, P. G., and Li, B. (2013). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, 1–23.
- Ghose, A., and Yang, S. (2009). An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science*, 55 (10).
- Godbole, N., Srinivasaiah, M., and Skiena, S. (2007). Large-scale sentiment analysis for news and blogs. In Proceedings of the international conference on weblogs and social media (icwsm).

- Hariharan, S., and Ramkumar, T. (2011). Mining product reviews in web forums. *Int. J. Inf. Retr. Res.*, 1 (2), 1–17.
- Hariharan, S., Srimathi, R., Sivasubramanian, M., and Pavithra, S. (2010). Opinion mining and summarization of reviews in web forums. In *Proceedings of the third annual acm bangalore conference* (pp. 241–244). New York, USA: ACM.
- Hatzivassiloglou, V., and Wiebe, J. M. (2000). Effects of adjective orientation and gradability on sentence subjectivity. In *Proceedings of the 18th conference on computational linguistics - volume 1* (pp. 299–305). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Hu, C., Weng, Y., Zhang, X., and Xue, C. (2009). Hot topic detection based on opinion analysis for web forums in distributed environment. In *Intelligent distributed computing iii* (Vol. 237, p. 101-110). Springer Berlin Heidelberg.
- Hu, M., and Liu, B. (2004a). Mining and summarizing customer reviews. In *Proceedings of the tenth acm sigkdd international conference on knowledge discovery and data mining* (pp. 168–177). New York, USA: ACM.
- Hu, M., and Liu, B. (2004b). Mining opinion features in customer reviews. In (pp. 755–760). AAAI Press.
- Huang, R., Sun, L., and Feng, Y. (2008). Study of kernel-based methods for Chinese relation extraction. In H. Li, T. Liu, W.-Y. Ma, T. Sakai, K.-F. Wong, and G. Zhou (Eds.), *Information retrieval technology* (Vol. 4993, p. 598-604). Springer Berlin Heidelberg.
- Hui, P., and Gregory, M. (2010). Quantifying sentiment and influence in blogspaces. In *Proceedings of the first workshop on social media analytics* (pp. 53–61). New York, USA: ACM.
- Jansen, B. J., Spink, A., and Pedersen, J. (2005). A temporal comparison of altavista web searching: Research articles. *J. Am. Soc. Inf. Sci. Technol.*, 56 (6), 559–570.
- Jiang, L., Yu, M., Zhou, M., Liu, X., and Zhao, T. (2011). Target-dependent twitter sentiment classification. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies - volume 1* (pp. 151–160). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Joshi, M., and Penstein-Rosé, C. (2009). Generalizing dependency features for opinion mining. In *Proceedings of the acl-ijcnlp 2009 conference short papers* (pp. 313–316). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Khan, A., and Baharudin, B. (2011a). Sentiment classification by sentence level semantic orientation using sentiwordnet from online reviews and blogs. *International Journal of Computer Science & Emerging Technologies*, 2 (4), 539–552.
- Khan, A., and Baharudin, B. (2011b). Sentiment classification using sentence-level semantic orientation of opinion terms from blogs. In *National postgraduate conference (npc)* (p. 1-7).
- Kim, J.-D., Ohta, T., Pyysalo, S., Kano, Y., and Tsujii, J. (2009). Overview of bionlp'09 shared task on event extraction. In *Proceedings of the workshop on current trends in biomedical natural language processing: Shared task* (pp. 1–9). Stroudsburg, PA, USA: Association for Computational Linguistics.

- Kobayashi, N., Iida, R., Inui, K., and Matsumoto, Y. (2005). Opinion extraction using a learning-based anaphora resolution technique. In Proceedings of the 2nd international joint conference on natural language processing (ijcnlp- 05), poster (p. 175-180).
- Kong, R., Wang, Y., Xin, W., Yang, T., Hu, J., and Chen, Z. (2011). Customer reviews for individual product feature-based ranking. In Proceedings of the 2011 first international conference on instrumentation, measurement, computer, communication and control (pp. 449–453). IEEE Computer Society.
- Levy, R., and Manning, C. (2003). Is it harder to parse Chinese, or the Chinese treebank? In Proceedings of the 41st annual meeting on association for computational linguistics - volume 1 (pp. 439–446).
- Li, C.-R., Yu, C.-H., and Chen, H.-H. (2011). Predicting the semantic orientation of terms in e-hownet. In Proceedings of the 23rd conference on computational linguistics and speech processing (pp. 151–165). Stroudsburg, PA, USA.
- Li, J., and Sun, M. (2007, Aug). Experimental study on sentiment classification of Chinese review using machine learning techniques. In International conference on natural language processing and knowledge engineering (pp. 393–400).
- Liu, B. (2006). Web data mining: Exploring hyperlinks, contents, and usage data (data-centric systems and applications). Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Liu, B. (2011). Web data mining: Exploring hyperlinks, contents, and usage data (data-centric systems and applications). Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Liu, B. (2012). Sentiment analysis and opinion mining. Synthesis Lectures on Human Language Technologies, 5 (1), 1–167.
- Liu, B., Hu, M., and Cheng, J. (2005). Opinion observer: Analyzing and comparing opinions on the web. In Proceedings of the 14th international conference on world wide web (pp. 342–351). New York, USA: ACM.
- Marchetti-Bowick, M., and Chambers, N. (2012). Learning for microblogs with distant supervision: Political forecasting with twitter. In Proceedings of the 13th conference of the european chapter of the association for computational linguistics (pp. 603–612). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Matsumoto, S., Takamura, H., and Okumura, M. (2005). Sentiment classification using word sub-sequences and dependency sub-trees. In Proceedings of the 9th pacific-asia conference on advances in knowledge discovery and data mining (pp. 301–311). Berlin, Heidelberg: Springer-Verlag.
- Meena, A., and Prabhakar, T. V. (2007). Sentence level sentiment analysis in the presence of conjuncts using linguistic analysis. In Proceedings of the 29th european conference on ir research (pp. 573–580). Berlin, Heidelberg: Springer-Verlag.
- Melville, P., Gryc, W., and Lawrence, R. D. (2009). Sentiment analysis of blogs by combining lexical knowledge with text classification. In Proceedings of the 15th acm sigkdd international conference on knowledge discovery and data mining (pp. 1275–1284). New York, USA: ACM.
- Mihalcea, R., Banea, C., and Wiebe, J. (2007, June). Learning multilingual subjective language via cross-lingual projections. In Proceedings of the association for computational linguistics (acl) (pp. 976–983). Prague, Czech Republic: Association for Computational Linguistics.

- Montejo-R´aez, A., Mart´inez-C´amara, E., Mart´ın-Valdivia, M. T., and Uren˜a L´opez, L. A. (2014). Ranked wordnet graph for sentiment polarity classification in twitter. *Comput. Speech Lang.*, 28 (1), 93–107.
- Moraes, R., Valiati, J. a. F., and Gavi˜ao Neto, W. P. (2013). Document-level sentiment classification: An empirical comparison between svm and ann. *Expert Syst. Appl.*, 40 (2), 621–633.
- Morinaga, S., Yamanishi, K., Tateishi, K., and Fukushima, T. (2002). Mining product reputations on the web. In *Proceedings of the eighth acm sigkdd international conference on knowledge discovery and data mining* (pp. 341–349). New York, USA: ACM.
- Mukherjee, S., and Bhattacharyya, P. (2012). Wikisent: Weakly supervised sentiment analysis through extractive summarization with Wikipedia. In *Machine learning and knowledge discovery in databases* (Vol. 7523, pp. 774–793). Springer Berlin Heidelberg.
- Nasukawa, T., and Yi, J. (2003). Sentiment analysis: Capturing favorability using natural language processing. In *Proceedings of the 2nd international conference on knowledge capture* (pp. 70–77). New York, USA: ACM.
- Pandarachalil, R., and Sendhilkumar, S. (2013). Identifying same wavelength groups from twitter: A sentiment based approach. In *Proceedings of the 5th Asian conference on intelligent information and database systems - volume part ii* (pp. 70–77). Berlin, Heidelberg: Springer-Verlag.
- Pang, B., and Lee, L. (2004). A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd annual meeting on association for computational linguistics*. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Pang, B., Lee, L., and Vaithyanathan, S. (2002). Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the acl-02 conference on empirical methods in natural language processing - volume 10* (pp. 79–86). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Popescu, A.-M., and Etzioni, O. (2005). Extracting product features and opinions from reviews. In *Proceedings of the conference on human language technology and empirical methods in natural language processing* (pp. 339–346). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Pu, P., Chen, L., and Kumar, P. (2008). Evaluating product search and recommender systems for e-commerce environments. *Electronic Commerce Research*, 8 (1-2), 1–27.
- Ravi Kumar, V., and Raghuvver, K. (2013). Dependency driven semantic approach to product features extraction and summarization using customer reviews. In *Advances in computing and information technology* (Vol. 178, pp. 225–238).
- Saranya, T. (2013). Mining features and ranking products from online customer reviews. *International Journal of Engineering Research & Technology (I- JERT)*, 2 (10), 643–648.
- Silverstein, C., Marais, H., Henzinger, M., and Moricz, M. (1999). Analysis of a very large web search engine query log. *SIGIR Forum*, 33 (1), 6–12.
- Souza, M., and Vieira, R. (2012). Sentiment analysis on twitter data for Portuguese language. In *Proceedings of the 10th international conference on computational processing of the Portuguese language* (pp. 241–247). Berlin, Heidelberg: Springer-Verlag.

- Speriosu, M., Sudan, N., Upadhyay, S., and Baldrige, J. (2011). Twitter polarity classification with label propagation over lexical links and the follower graph. In Proceedings of the first workshop on unsupervised learning in nlp (pp. 53–63). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Su, Y., and Li, S. (2012). Constructing Chinese sentiment lexicon using bilingual information. In Proceedings of the 13th Chinese conference on Chinese lexical semantics (pp. 322–331). Berlin, Heidelberg: Springer-Verlag.
- Subrahmanian, V. S., and Reforgiato, D. (2008, July). Ava: Adjective-verb- adverb combinations for sentiment analysis. *Intelligent Systems, IEEE* , 23 (4), 43-50.
- Tan, C., Lee, L., Tang, J., Jiang, L., Zhou, M., and Li, P. (2011). User-level sentiment analysis incorporating social networks. In Proceedings of the 17th acm sigkdd international conference on knowledge discovery and data mining (pp. 1397–1405). New York, USA: ACM.
- Tong, R. (2001). An operational system for detecting and tracking opinions in on-line discussions. In Working notes of the sigir workshop on operational text classification (pp. 1–6). New Orleans, Louisiana.
- Tsou, B. K., Yuen, R. W., Kwong, O., and Lai, W. L., Tom B.Y.and Wong. (2005). Polarity classification of celebrity coverage in the Chinese press. In Proceedings of international conference on intelligence analysis. McLean, VA, USA.
- Turney, P. D. (2002). Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. In Proceedings of the 40th annual meeting on association for computational linguistics (pp. 417–424). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Wan, X. (2008). Using bilingual knowledge and ensemble techniques for unsupervised Chinese sentiment analysis. In Proceedings of the conference on empirical methods in natural language processing (pp. 553–561). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Wang, D., Zhu, S., and Li, T. (2013). Sumview: A web-based engine for summarizing product reviews and customer opinions. *Expert Systems with Applications*, 40 (1), 27 – 33.
- Wang, S., Wei, Y., Li, D., Zhang, W., and Li, W. (2007). A hybrid method of feature selection for Chinese text sentiment classification. In Proceedings of the fourth international conference on fuzzy systems and knowledge discovery - volume 03 (pp. 435–439). Washington, DC, USA: IEEE Computer Society.
- Wiebe, J. (2000). Learning subjective adjectives from corpora. In Proceedings of the seventeenth national conference on artificial intelligence and twelfth conference on innovative applications of artificial intelligence (pp. 735– 740). AAAI Press.
- Wiebe, J., Wilson, T., Bruce, R., Bell, M., and Martin, M. (2004). Learning subjective language. *Comput. Linguist.*, 30 (3), 277–308.
- Wilson, T., Wiebe, J., and Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In Proceedings of the conference on human language technology and empirical methods in natural language processing (pp. 347–354). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Wu, F., and Weld, D. S. (2010). Open information extraction using Wikipedia. In Proceedings of the 48th annual meeting of the association for computational linguistics (pp. 118–127). Stroudsburg, PA, USA: Association for Computational Linguistics.

- Wu, Y., Zhang, Q., Huang, X., and Wu, L. (2009). Phrase dependency parsing for opinion mining. In Proceedings of the 2009 conference on empirical methods in natural language processing: Volume 3 - volume 3 (pp. 1533–1541). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Wu, Y., Zhang, Q., Huang, X., and Wu, L. (2011). Structural opinion mining for graph-based sentiment representation. In Proceedings of the conference on empirical methods in natural language processing (pp. 1332–1341). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Xiang, G., Fan, B., Wang, L., Hong, J., and Rose, C. (2012). Detecting offensive tweets via topical feature discovery over a large scale twitter corpus. In Proceedings of the 21st acm international conference on information and knowledge management (pp. 1980–1984). New York, USA: ACM.
- Ye, Q., Shi, W., and Li, Y. (2006). Sentiment classification for movie reviews in Chinese by improved semantic oriented approach. In Proceedings of the 39th annual Hawaii international conference on system sciences - volume 03 (pp. 3–5). Washington, DC, USA: IEEE Computer Society.
- Yu, J., Zha, Z.-J., Wang, M., and Chua, T.-S. (2011). Aspect ranking: Identifying important product aspects from online consumer reviews. In Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies - volume 1 (pp. 1496–1505). Association for Computational Linguistics.
- Zhang, K., Cheng, Y., Liao, W.-k., and Choudhary, A. (2012). Mining millions of reviews: A technique to rank products based on importance of reviews. In (pp. 121–128). New York, USA: ACM.
- Zhang, K., Narayanan, R., and Choudhary, A. (2010). Voice of the customers: Mining online customer reviews for product feature-based ranking. In (pp.11–11). Berkeley, CA, USA: USENIX Association.
- Zhang, X., and Zhou, Y. (2011). Holistic approaches to identifying the sentiment of blogs using opinion words. In Proceedings of the 12th international conference on web information system engineering (pp. 15–28). Berlin, Heidelberg: Springer-Verlag.
- Zhu, Y.-l., Min, J., Zhou, Y.-q., Huang, X.-j., and Li-de, W. (2006). Semantic orientation computing based on hownet. *Journal of Chinese Information Processing*, 14–20.
- Zhuang, L., Jing, F., and Zhu, X.-Y. (2006). Movie review mining and summarization. In Proceedings of the 15th acm international conference on information and knowledge management (pp. 43–50). New York, USA: ACM.