

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

A Review of Opinion Mining and Sentiment Classification Framework in Social Networks

Yee W. LO¹, DEBI and Vidyasagar POTDAR², DEBI

¹Yee W. LO, DEBI Institute, Curtin University of Technology, Australia, e-mail : yeewen.lo@postgrad.curtin.edu.au

²Vidyasagar. POTDAR, DEBI Institute, WA, Australia
e-mail: v.potdar@curtin.edu.au

Abstract—The Web has dramatically changed the way we express opinions on certain products that we have purchased and used, or for services that we have received in the various industries. Opinions and reviews can be easily posted on the Web, such as in merchant sites, review portals, blogs, Internet forums, and much more. These data are commonly referred to as user-generated content or user-generated media. Both the product manufacturers, as well as potential customers are very interested in this online ‘word-of-mouth’, as it provides product manufacturers information on their customers likes and dislikes, as well as the positive and negative comments on their products whenever available, giving them better knowledge of their products limitations and advantages over competitors; and also providing potential customers with useful and ‘first-hand’ information on the products and/or services to aid in their purchase decision making process. This paper discusses the existing works on opinion mining and sentiment classification of customer feedback and reviews online, and evaluates the different techniques used for the process. It focuses on the areas covered by the evaluated papers, points out the areas that are well covered by many researchers and areas that are neglected in opinion mining and sentiment classification which are open for future research opportunity.

Index Terms—Opinion Mining, Product Reviews, Market Intelligence, Customer Reviews.

I. INTRODUCTION

As mentioned in many literatures and reports, the Web contains a wealth of product reviews by customers due to the ease of publishing online in the recent years. This leads to an explosion in the web content of customer reviews and opinions on products. It is an expensive and daunting task to sift through them [2]. The good news is that this mass of data makes automation of online customer reviews collection and classification possible and worthwhile.

The mining, classification and collection of information from these online customers are important to the product manufacturers, as it provides them information about any product defects at an early stage from customers complaints that could quickly proliferate through the Internet. Knowing products limitations and defects can be helpful in risk management and henceforth, reducing future liabilities, as well as to make sound marketing strategies. These online reviews are also interesting to existing and potential customers too. It could significantly influence their purchase decisions of a product or in hiring a service. These explain for the importance of opinion mining and sentiment classification of online customer reviews.

To further explain, the goal of business intelligence in

customer satisfaction studies of their opinions on products and/or services, and can often be explored using surveys and focus groups studies. However, these can incur high expenses associated with the design and the administration of the surveys, and also are time consuming in the process of the obtainment of the resulting data [10].

A less formal or structured alternative is the collection of spontaneous customer reviews and feedbacks on the Web, such as blogs, newsgroups, feedback emails and review websites. These data can come in free text that is less structured compared to surveys and focus groups. However, without much effort, free-form customer opinions can also be processed more efficiently and effectively, using automatic text mining techniques in the market, such as clustering and key term extraction [10].

The recent swell of interest in the automatic identification and extraction of attitudes, opinions, and sentiments in text, is due to researchers’ desire to provide systems that can help and support information analysts in the various industries, with the mundane and time consuming task of sifting through the massive amounts of web data, trying to answer questions like: ‘How do people feel about the latest camera phone?’ or ‘Is there a change in the support for the new Medicare bill? [4]’

This is when Opinion Mining (OM – also known as “sentiment classification”) comes in. OM is a recent subdiscipline at the crossroads of information retrieval and computational linguistics that is concerned with the opinion a document expresses [11]. The task of Opinion Mining can be classified into two subtasks: Document classification and information extraction [5]. Document classification is the classifying of documents or passages according to semantic/sentiment orientation such as positive vs. negative [5]; and information extraction focuses on the extraction of opinions consisting of information about particular aspects of interest and the corresponding sentiment orientation in a structured form from a set of unstructured data, which is in contrast to sentiment classification in general where the former aims at producing richer information useful for in-depth analysis of opinions, which has recently been challenged by the growing research community [5].

This research paper is structured into various sections. Section II is the literature review that presents the existing and related work on Opinion Mining, depicting the major contributions and limitations of the existing techniques. Section III follows with the evaluation of the researches. Our recommendations for future research opportunity are

reported in the conclusion in Section IV.

II. LITERATURE REVIEW

In this section, we present a review of the existing and related works on Opinion Mining (OM) and Sentiment Classification (SC) proposed in the literature. The state-of-the-art theories and models in today's literature are also presented. This review is categorized in the following seven categories as shown in Figure 1. It outlines the various techniques used for Opinion Mining and Sentiment Classification from the existing literature. The different techniques used to mine opinions, classify sentiment of mined items and features, as well as the strength of the sentiment are reviewed; and compared and contrasted against each other.

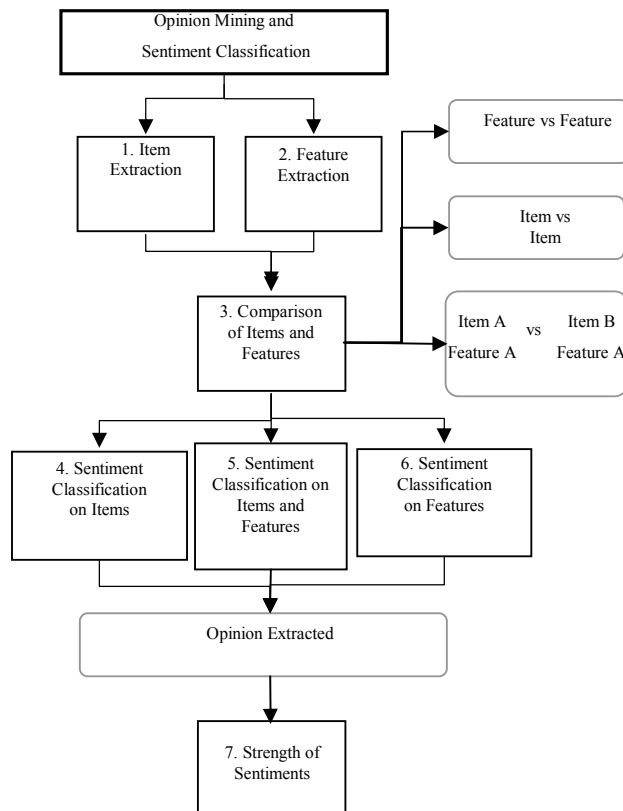


Fig 1. Flowchart of OM and SC Framework

A. Item Extraction

In this section, we analyze all the frameworks that are related to *item extraction*. Specifically two papers are of importance as they focus on this topic in detail.

Item extraction is the process of extracting the subject matter where opinions have been expressed on in customer reviews. It is also commonly referred to as ‘Opinion Extraction’ [5]. The term ‘subject matter’ is also commonly referred to as a situation or a product. It is an important task as it is the beginning stage in the task of OM. In contrast to sentiment classification, opinion extraction aims at producing useful richer information for in-depth analysis of opinions [5].

When evaluating these frameworks, we rely on the questions listed in Table 1. Answers to these questions reflect the advantages and disadvantages of the researched frameworks.

Table 1. An evaluation of the Item Extraction methods		
Papers reviewed	[5]	[10]
Parameters for evaluation		
Is machine-learning technique being applied in the algorithm?	Y	Y
Does the algorithm uses supervised learning?	Y	N
Is algorithm portable across domain?	Y	N
Does the algorithm extracts and combines contextual clues and context independent statistical clues?	Y	N
Does the algorithm uses a combination of clustering techniques and machine-learned sentiment classifier?	N	Y
Does the algorithm shows significant improvement over existing methods?	Y	Y

Since this problem was discovered, limited research has been undertaken to attempt to solve it, for example the work by Kobayashi et al. [5] and Gamon et al. [10].

Kobayashi and his team presented a method for opinion extraction in a structured form. It also discussed the most effective way to structure customer reviews in web documents and focused on extracting subject/aspect-evaluation relations, and extracting subject/aspect-aspect relations, using a machine learning-based method, which is portable across domains. It addressed the task of opinion extraction by combining contextual clues and context-independent statistical clues using a machine-learning technique, ‘boosting-based algorithm’ by Kudo (04) implemented as the package BACT. Experiments were carried out and evaluation was conducted using 5 fold cross validation on all data in the aspects of recall and precision.

Compared the Kobayashi [5], the approach presented by Gamon and his team [10], is termed as *Pulse*, a prototype system for mining topics and sentiment orientation jointly from free text customer feedback. *Pulse* combines a clustering technique with a machine-learned sentiment classifier, allowing for a visualization of topic and associated customer sentiment. The Tree Map visualization is used to display clusters and their associated sentiment. It allowed the identification of the overall sentiment associated with product make/model, the most common topics that customers mentioned in reviews, as well as the most positive and the most negative topics, at a glance.

Table 1 shows that machine-learning techniques have been applied in both algorithms and experimental results have shown significant improvement over many existing methods such as Baseline A-E model and Context-only A-E model; where Baseline A-E (aspect-evaluation) model stimulates the algorithm proposed by (Tateishi 04) and Context-only A-E model uses contextual pattern-based clues but not statistical clues and is a boost-based algorithm proposed by Kudo (04).

Table 1 also shows that Kobayashi [5] used supervised learning and a combination of contextual clues and context-

independent statistical clues in their experiment; whilst Gamon [10] used a combination of clustering techniques and machine-learned sentiment classifier.

We will explore another area of opinion mining, ‘feature extraction’, in the next section that can add more value to the existing knowledge about the item extracted from reviews.

B. Feature Extraction

Feature extraction is the identification of features of products which customers have expressed their opinions on their reviews and feedbacks. Features refer to product features, product attributes, and/or product functions like the picture quality of the Canon IXUS 10, or the interior design of a Ford territory, or the service of hotel staff.

It is essential to readers that the features of the reviewed products are known as their areas of importance in different products may differ from people. For example, a reader might be more interested in the cleanliness of the hotel room, whilst the reviewer is more concerned with the quality of the customer service of the hotel staff.

Having said, current work in the domain of Feature Extraction is still in its infancy. We have reviewed three papers and identified their main contributions, weaknesses, similarities and differences as shown in Table 2.

Table 2. An evaluation of Feature Extraction Methods			
Papers Reviewed	1	6	8
Parameters for evaluation			
Does the algorithm identify & classify opinion sentences in reviews?	Y	N	N
Does the algorithm provide a summarization of the results?	Y	N	N
Does the algorithm use POS Tagging?	Y	N	Y
Does the algorithm use WordNet?	Y	N	Y
Does the algorithm use Word n-grams?	N	N	Y
Can algorithm detect the weight of the opinion based on the opinion itself?	N	N	N
Can algorithm differentiate different features of the same product?	Y	Y	Y
Can feature hierarchy be constructed?	Y	N	N

Hu et al. [1] studied the problem of generating feature-based summaries (FBS) of customer reviews of products sold online and strived to mine and summarize all the customer reviews. Their task was performed in three main steps. Firstly, to mine product features that have been commented on by customers; secondly, to identify opinion sentences in each review and decide whether each opinion sentence is positive or negative; and finally to summarize the results.

The paper also proposed several novel techniques to aid in the process of performing these tasks such as POS Tagging, Association Miner CBA, and WordNet. Their experimental results using a large number of customer reviews of 5 products sold online demonstrated the effectiveness of the Feature-Based Summarization (FBS)

and its techniques against other existing methods used by other researchers such as the FASTR of Christian Jacquemin [14].

In this work, opinion classification was performed at the sentence level rather than at the document level. In contrast to related works like Dave, Lawrence and Pennock’s work [2], they did not find features on which opinions have been expressed, or summarized similarities and differences of reviews, but instead aimed to find the key features that are talked about in multiple reviews. Their focus was on using association mining to find all frequent features.

The next paper by Kobayashi et al. [6] not only extracted the explicit features in reviews, they also extracted the *Subject* and the *Value* of the reviews. They believed these details are essential for addressing the task.

Assuming opinions can be represented as tuples (Subject, Attribute, Value), they addressed the task of opinions extraction by employing a computational method for tuples extraction. Machine-learning based techniques are then applied to the main task of opinion extraction, which was then decomposed into two subtasks: Extraction of attribute-value pairs related to a product (where an attribute represents one aspect of a subject and the value is a specific language expression that qualifies or quantifies the aspect); and Determination of its subjectivity on the opinion as a whole. The proposed method had yield a better outcome.

The next paper to be discussed is from Mishne [8] who took a totally different and unique approach, the first in the database of literature on opinion mining at that time. Mishne’s work as was more focused on the classification of the blog posts by different moods. In other words, the paper not only addressed the task of feature extraction, but continued to use the extracted features in assisting for their task of classifying blog posts by moods. They believed that mood classification is useful for various applications, such as assisting behavioral scientists and improving doctor-patient interaction. Their objective was to predict the reviewer’s most likely state of mind when the post was written using a machine learning approach to identify a set of features to be used for the learning process.

Their experimental results had showed a consistent modest improvement on the naïve baseline.

Table 2 shows that Hu [1] and Mishne [8] had used POS Tagging and Word n-grams in their work. Hu [1] also used Association Miner CBA, and WordNet; whilst Mishne used Naïve method, frequencies of word lemmas acquired from Tree Tagger [16], as well as SVMlight package in their work. Contrasted, Kobayashi [6] applied a machine learning-based method used for anaphora resolution to the opinion extraction problem. Table 2 also shows that all the three algorithms can differentiate different features of the same product, but only Hu’s work [1] allowed feature hierarchy to be constructed as they focused on finding key features of the products.

The tasks of item and feature extraction are covered in section A and B respectively. In section C, the task of sentiment classification is discussed.

C. Sentiment Classification in General

Sentiment classification is the process of determining the subjectivity of a given text. In simple words, it is the task of deciding whether a given text expresses a positive or negative opinion about its ‘subject matter’ and ‘subject attributes’, which is also known as ‘product’ and ‘features’.

Papers reviewed	2	7	11	12
Parameters for evaluation				
Does the algorithm uses scoring methods from information retrieval for sentiment determination?	Y	N	N	N
Does the algorithm use General Inquirer (GI) lexicon?	N	N	Y	Y
Does the algorithm uses WordNet?	N	N	N	Y
Does the algorithm use linguistic rules?	N	Y	N	N
Does the algorithm use aggregation function?	N	Y	N	N

Sentiment classification assessment on opinions is usually done on document level rather than on sentence level as in contrast to opinion mining [5].

Table 3 presents the similarities and differences of the four papers relevant to the area of sentiment classification.

Dave et al. [2] applied various machine learning methods for the task of opinion extraction and sentiment classification and discovered several problems that have not been expected initially.

They began by using structured reviews for testing and training and identifying appropriate features. Table 2 shows that scoring methods was used in information retrieval for determination of the sentiment of reviews.

Experiments were conducted on user reviews on C|net and Amazon where authors provided quantitative or binary ratings, which were believed to be perfect for training and testing for sentiment orientation.

Two tests were conducted and experimental results showed that their best methods performed as well as or better than traditional machine learning methods.

Sentiment classification in general is covered and the next section is ‘sentiment classifications on items’ which provides more specified and targeted information.

D. Sentiment Classification on Item

This section is more focused on whether a given text has positive or negative connotation on its subject matter only. For example, Camera 1 has positive or negative feedback from users online.

The paper by Esuli et al. [11] is the extension of the authors’ previous work: “Esuli & Sebastiani, 2005, ‘Determining the semantic orientation of terms through gloss analysis’ [13]”.

It confronted the task on the decision of whether a given term has a positive connotation, or a negative connotation, or has no subjective connotation at all; thus, this problem subsumed the problem of determining orientation. This problem was tackled by testing three different variants of a

semi-supervised method previously proposed for orientation detection. Their results showed that determining subjectivity orientation was much of a harder problem than determining orientation alone.

The benchmark that they had used for their experiments is the General Inquirer (GI) lexicon (Stone et al., 1996) as shown in Table 3. This is a lexicon of terms labeled according to a large set of categories, each one denoting the presence of a specific trait in the term.

Unfortunately, their results had shown that an algorithm that had shown excellent state-of-the-art performance in deciding term orientation [13], once modified for the purposes of deciding term subjectivity, performed more poorly. This had been shown by testing several variants of the basic algorithm, some of them involving radically different supervised learning policies.

The next work by Esuli et al. [12] used *SentiWordNet*, a lexical resource describing the degree of positivity and negativity of the extracted terms and features. It is not an extension of the previous research, although Table 3 has showed that both used GI lexicon.

SENTIWORDNET, is a lexical resource in which each WORDNET synset, *s*, is associated to three numerical scores Obj (*s*), Pos (*s*) and Neg (*s*), describing how objective, positive, and negative the terms contained in the synset are.

Based on observations in their previous works, they had decided to combine different configurations of training set and learner into a committee to produce the final SENTIWORDNET scores. They believed SentiWordNet is a useful tool in opinion mining applications, because of its wide coverage and its fine grain properties, obtained by qualifying the labels by means of numerical scores.

The next section is sentiment classification on features, which is the next stage in sentiment classification on items.

E. Sentiment Classification on Features

We are discussing sentiments on features in this section, as in deciding whether a given text has a Positive or Negative opinion on its ‘subject attributes’, which we also commonly referred to as ‘product attributes’ and/or ‘product features’. For example, the features of Camera A being ‘buttons placement on the camera’ and ‘size of screen’.

Of particular relevance in this topic is the work by Ding et al.’s [7]. Ding [7] discussed the problem of determining semantics of opinions expressed on product features in customer reviews, rather than on the products (items) mentioned in the reviews compared to Esuli et al. in [11] and [12].

The objective of Ding’s work was to use linguistic rules together with a new opinion aggregation function to address the problem of sentiment classification on features as shown in Table 3. This approach used context to infer the orientations of opinions on a product feature. The ‘S system’, called Opinion Observer, was also presented in this research.

In this work, the product features are assumed to be given or discovered before determining whether an opinion

is positive or negative and results of experiments on Opinion Observer showed that it had outperformed both FBS [1] and OPINE [3]. Additional experiments had also showed that both the new opinion aggregation function and linguistic rules had contributed roughly equally to the improved results of Opinion Observer over FBS with the recall factor improving dramatically without much loss in the precision rating.

We have covered the different aspects of sentiment classifications in section C, D and E, as well as the different methods used to accomplish the task effectively.

In the next section, we touched on the ‘strengths of the sentiments’ to get more detailed information and a better understanding on the extracted sentiments.

F. Strength of Sentiments

Determining the strength of sentiments is the process of deciding whether a Positive opinion expressed by a text on its subject matter is Weakly Positive, Mildly Positive, or Strongly Positive, and/or whether a Negative opinion expressed is Weakly Negative, Mildly Negative, or Strongly Negative.

Parameters for evaluation	Papers reviewed	
	3	4
Is algorithm portable across domain (Domain independent)?	Y	N
Does the algorithm categorize the sentiment strength using objective measures?	Y	Y

In this section, the two papers of particular relevance are Popescu et al. [3] and Wilson & Wiebe et al. [4], whose works discuss the strength of the sentiments of customer reviews and feedbacks.

The paper by Popescu et al. [3] introduced OPINE, a more in-depth and detailed unsupervised information extraction system which extracted fine-grained features, and associated opinions, from reviews. OPINE mined reviews for the purpose of building a model of important product features, their evaluation by reviewers, and their relative quality across products. It focused on the extraction of explicit features, identifying corresponding customer opinions about these features and determining their polarity. It differed from method used in [1] as instead of only finding the key features in multiple reviews, it also determines the polarity and strengths of the sentiments of the reviews. It is also portable across domain as shown in Table 4.

OPINE extracts product features such as properties, parts, features of product parts, related concepts, and parts and properties of related concept, and also opinion phrases, which are adjective, noun, verb or adverb phrases representing customer opinions using WordNet’s IS-A hierarchy and morphological cues from a set of reviews from Hu & Liu’s [1] publicly available data sets.

After various tests and experiments, results had confirmed the performance superiority of OPINE in its tasks

as compared to similar review-mining system, FBS [1]. Its novel use of relaxation labeling technique for determination of semantic orientation of words in context of given product features and sentences had also led to better performance on the tasks of customer opinions identification and their polarity. OPINE was literally an improvisation of existing Hu’s FBS on feature extraction [1].

The work by Wilson and Wiebe et al. [4] was the first research that touches on automated opinion and sentiment classification. Its main objective was to automate the task of distinguishing between objective and subjective languages as well as the strength of the sentiments of each sentence in the corpus of reviews used for their paper, which was very different to OPINE by Popescu et al. [3].

Experiments were conducted and they had achieved significant improvements in mean-squared error over baseline using three machine learning algorithms, which are BoosTexter (Schapire and Singer, 2000) AdaBoost.HM for boosting, Ripper (Cohen, 1995) for rule learning, and SVMlight (Joachims, 1999) for support vector regression.

In the next section, the comparison of items and features is discussed.

G. Comparison of Items and Features

The comparison of items and features is the process of comparing products in the same product groups (e.g. Canon camera and Sony camera) mentioned in customer reviews in terms of their features. E.g. “the battery life of camera A is much longer than the battery life of camera B”, or “I like the lens of camera A but at the same time I prefer the battery life of Camera B”. The purpose of the comparison is to allow readers to quickly and clearly identify the more superior product out of many in terms of different features.

The work by Liu et al. [9] addressed this topic by proposing an analysis system, *Opinion Observer*, with a visual component for comparison of consumer opinions on different products. This system compared and contrasted features of various products in many forms of visual diagrams allowing users to clearly see the strengths and weaknesses of products based on features at a single glance.

Their work was related but quite different from sentiment classification; as its purpose was to classify reviews as positive or negative and did not identify product features. Experimental results on the evaluation of their proposed technique in identification of product features from Pros and Cons had shown that the technique was highly effective and outperformed existing methods significantly.

We can see that the Opinion Observer is a very powerful and useful analysis system created by Liu et al. to allow readers to quickly absorbed information in numerous reviews at a single glance to make informed purchase decision. Without it, readers will have to spend many hours going through mountainous of online reviews which are tedious and unconventional in our today’s time-pressed society.

III. EVALUATION & DISCUSSION

This section of the paper uncovers areas that have not been studied much in the existing literature. As illustrated in Fig. 2, most of the research so far is focused on the task of feature extraction and sentiment classification on features. Little work has been done on item extraction, sentiment classification of reviews, sentiment classification on items, and comparison of items and features.

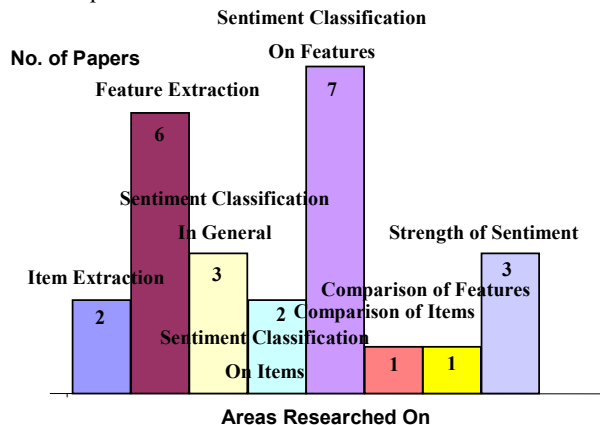


Fig. 2 Areas of OM the reviewed papers focused on

We believe that the comparison of items and features of products are important and significant in opinion mining of online customer reviews for the benefit of consumers, product distributors and product manufacturers. Having a visual chart or table which depicts the advantages and disadvantages of competing products and services at a glance is very useful. It allows the advantages and disadvantages of a product to its direct competing products to be quickly known, instead of having to spend huge amount of time going through infinite reviews on various websites. It also provides convenience and is less time-consuming for product manufacturers to gather market intelligence and product benchmarking information; and provides consumers with a side-by-side and feature-by-feature comparisons of consumer opinions on competing products [9]. The visuals allow potential customers to quickly figure out the information that they need before making a purchase decision, and relieves them from the mundane task of going through numerous reviews, especially in today's quick paced and time-pressed society where time is money.

An example can be shown in this section in Figure 2. It is obviously quicker and more convenient to pick up areas that are most researched on and areas that are least researched on from the visuals provided in Fig. 2 compared to reading the illustration in words.

IV. CONCLUSION

As seen in section II and III, many Opinion Mining methods, supervised and unsupervised, machine learning and non-machine learning, automated or manual, are discussed; as well as Sentiment Classification methods researched on and created by other researchers in the recent years. The possibility of employing various methods and

combining them to be used in a single situation to achieve better and more accurate outcome and results are also demonstrated in the respective sections.

We have also identified that the area, "Comparison on Items and Features" are the least researched on in the current literature on opinion mining and sentiment classification academic database. Thus, we hope this can be the future direction of research work for ourselves as well as other researchers.

V. REFERENCES

- [1] Hu, M. & Liu, B., 'Mining and Summarizing Customer Reviews', in *Proceedings of the 10th International Conference on Knowledge Discovery and Data Mining*, 2004, KDD-04, Seattle, WA, pp. 168-177.
- [2] Dave, K., Lawrence, S. & Pennock, D.M., 'Mining the Peanut Gallery: Opinion extraction and Semantic Classification of Product reviews', in *Proceedings of the 13th International World Wide Web Conference*, 2003, WWW03.
- [3] Popsecu, A.M. & Etzioni, O., 'Extracting Product Features and Opinions from Reviews', in *Proceedings of the Human Language Technology Conference (HLT) and Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2005, HLT/EMNLP-05, Vancouver, pp. 339-346.
- [4] Wilson, T., Wiehe, J. & Hwa, R., 'Just how mad are you? Finding strong and weak opinion clauses', in *Proceedings of the AAAI-04, 21st Conference of the American Association for Artificial Intelligence*, 2004, AAAI-04, San Jose, US, pp. 761-769.
- [5] Kobayashi, N., Inui, K. & Matsumoto, Y., 'Opinion Mining from Web Documents: Extraction and Structurization', in *Proceedings of the Transactions of the Japanese Society for Artificial Intelligence 22*, JSAI07, 2007, pp. 326-337.
- [6] Kobayashi, N., Iida, R., Inui, K. & Matsumoto, Y., 'Opinion Mining as Extraction of Attribute-Value Relations', in *Proceedings of the Nara Institute of Science and Technology*, JSAI 2005, 2005, Takayama, Ikoma, Japan, pp. 470-481.
- [7] Ding, X. & Liu, B., 'The Utility of Linguistic Rules in Opinion Mining', in *Proceedings of the SIGIR 2007*, SIGIR, 2007, Amsterdam, The Netherlands.
- [8] Mishne, G., 'Experiments with Mood Classification in Blog Posts', in *Proceedings of the Stylistic Analysis of Text for Information Access*, Style, 2005, Amsterdam, The Netherlands.
- [9] Liu, B., Hu, M. & Cheng, J., 'Opinion Observer: Analyzing and Comparing Opinions on the Web', in *Proceedings of the International World Wide Web Conference Committee*, WWW05, 2005, Chiba, Japan.
- [10] Gamon, M., Aue, A., Corston-Oliver, S. & Ringger, E., 'Pulse: Mining Customer Opinions from Free Text', in *Proceedings of the Natural Language Processing, Microsoft Research*, IDA, 2005, Redmond, WA, pp. 121-132.
- [11] Esuli, A. & Sebastiani, F., 'Determining Term Subjectivity and Term Orientation for Opinion Mining', in *Proceedings of the ACL-97, 35th Annual Meeting of the Association for Computational Linguistics*, CL-06, 2006, Madrid, ES, pp. 174-181.
- [12] Esuli, A. & Sebastiani, F., 'SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining', in *Proceedings of LREC 2006*, LREC, 2006.
- [13] Esuli, A. & Sebastiani, F., 'Determining the Semantic Orientation of Terms through Gloss Classification', in *Proceedings of 14th ACM International Conference on Information and Knowledge Management*, CIKM 05, 2005, Bremen, DE, pp. 617-624.
- [14] FASTER, <http://www.limsi.fr/Individu/jacquemi/FASTR/>
- [15] Pang, B. & Lee, L., 'Opinion mining and sentiment analysis', in *Proceedings of Foundation and Trends in Information Retrieval*, FTIR, 2008, Vol. 2, pp. 1-13
- [16] Schmid, H., 'Probabilities part-of-speech tagging using decision trees', in *Proceedings of International Conference on New Methods in Language Processing*, ICNMLP, 1994, Manchester, UK.