Marble Initiative

Monitoring the Impact of Events on Customers Opinion

M. Fernandes Caíña, R. Díaz Redondo and A. Fernández Vilas
I&C Lab. AtlantTIC Research Center, University of Vigo, 36310 Vigo, Spain

Keywords: Opinion Mining, Sentiment Analysis, Business Decision-Making, Longitudinal Analysis.

Abstract: Social networks have become a major source of information, opinions and sentiments about almost any subject. The purpose of this work is to provide evidences of the applicability of opinion mining methods to find out how some events may impact into public opinion about a brand, product or service. We report an experiment that mined Twitter data related to two particular brands during specific periods that have been selected from events that was supposed to affect the user's perception. To find out conclusions, the methodology of the experiment applies several pre-processing techniques to extract sentiment information from the posts (e.g., case alterations, Part-of-Speech tagging using a Natural Language Toolkit, symbols removal, sentence and n-gram separation). The *SenticNet 2 Corpus* is used for polarity classification by means of a supervised algorithm where several threshold values are defined to mark positive, negative and neutral opinions. A longitudinal inspection of the polarized results on histograms allows identifying the "hot spots" and relating them to real world events. Although this paper shows the finding in our initial experiments, the ultimate goal of the research initiative, which we call Marble, is to provide a cloud solution for early detection of opinion shifts by the automatic classification of events according to their impact on opinion (propagation speed, intensity and duration), and its relationship with the normal behavior around a brand, product or service.

1 INTRODUCTION

Internet has become more and more a communication and expression platform, rather than just a static information source. Mailing lists, forums and chats have been part of it since the very beginning, but over the last years, social networks have become the primary platform of communication for the majority of its users. *Facebook* and *Twitter* are the most notable examples, even if the latter is considered to be a microblogging platform rather than a social network.

For Twitter, the imposed limit of 140 characters encourages its users to post frequently without the associated hassle of writing an entire article in a blog site. This, along with a large user base, provides a network with a huge flow of real time information (about 500 million tweets per day in 2013 (Krikorian, 2013)) expressing opinions about almost everything, including events, experiences, products or services. Moreover, the public availability of the data turns this microblogging platform into an ideal vehicle to evaluate public opinion.

The value-added information offered by Twitter could provide an important insight about the real effect of certain business decisions on the customer's opinion, as well as environmental or external effects, and offers new indicators that could prove useful while managing the public image of a service, product or company. Decisions like maintaining a current line of marketing, retiring a product, selecting a damage control technique or continuing a viral campaign, could all benefit from this new data. Any company could use this new source of information on its benefit, improving or minimizing the impact on future business decisions.

The exchange of user's opinions throughout Internet is nothing new, as shopping and reviews sites have been collecting users opinions since more than a decade ago (e.g., Epinions, Amazon), but the main difference strikes in the method of expression at Social Media. Review sites usually give the users a customized form to fill, including pre-established categories for product features, rating fields, and some free form space for non-categorized information. On the contrary, Twitter lacks that

Fernandes Caíña M., Díaz Redondo R. and Fernández Vilas A..
Marble Initiative - Monitoring the Impact of Events on Customers Opinion.
DOI: 10.5220/0005151504030410
In *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval* (KDIR-2014), pages 403-410
ISBN: 978-989-758-048-2

structure and the information is exchanged in a totally free format.

Natural Language Processing (NLP), defined as the ability of a system to process human language (Preeti and BrahmaleenKaurSidhu, 2013), is an artificial intelligence component that can be used to mine opinion and sentiment from social networks, and classify each post as being positive, negative or neutral towards a specific subject, that is, to define a polarity for each one. According to (Pang and Lee, 2008), when broad interpretations are applied, “sentiment analysis” and “opinion mining” denote the same field of study (which itself can be considered a sub-area of subjectivity analysis). We use these terms more or less interchangeably.

By combining NLP and opinion mining and sentiment analysis techniques, a new research area emerges which allows processing free comments in Social Media and to infer the impact of business decisions such as a product reformulation or a new service offering, but also external phenomena such as new competitor’s’ strategies. For the experiment in this paper, two consumer electronics brands and two specific time intervals were selected with the objective of measuring the impact of external and internal events on the users’ opinion in Social Media (Twitter). After extracting the polarity of Twitter posts, events’ impact was evaluated and measured according to three features of the Social response: intensity, propagation speed and duration.

This first experiment is part of a research initiative, which we call *Marble*, a platform to assist decision making on the fly by continuously monitoring Twitter posts to (1) detect signs of opinion variations about brands and (2) discover causation from corporative internal information (so internal business decision) and from outside information in the Web (external context not controlled by the brand’s strategy). This kind of early assistance is essential today since user’s opinion is a direct indicator of the satisfaction associated with a company, but it also could affect the brand’s perception on their followers, with the potential risk of becoming viral.

2 RELATED WORK

NLP is on its own a big area of current research and development with a quite wide range of toolkits. *Apache OpenNLP* (The Apache Software Foundation, 2014), *Stanford CoreNLP* (The Stanford Natural Language Processing Group, 2014) and NLTK (Natural Language Toolkit for Python Project, 2014) represent some of the most important names in this development area which offers different types of classifiers, tokenizers and corpora to be customized according to the needs of the application purposes. Leaving apart language processing, numerous lexical databases like WordNet® (Miller, 1995) have been created to map the words functions and meanings. One step further, SentiWordNet (Esuli and Sebastiani, 2006) adds opinion polarity and affective information at a syntactical level to the WordNet® data, and is available to be used by opinion mining systems. SenticNet 2 corpus (Cambria et al., 2012) represents another lexical database that differentiates itself from SentiWordNet by including polarity and affective information for not only words but common sense knowledge concepts (e.g., phrases), commonly used to express an opinion.

Multiple models have been proposed to implement the whole opinion mining process. The work shown in (Pak and Paroubek, 2010) introduces a methodology to collect a corpus of Twitter posts to train a sentiment classifier. This classifier will be able to determine the polarity of a text using a multinomial Naïve Bayes classifier. Also with Twitter as workbench, in (Gokulakrishnan et al., 2012) some preprocessing techniques (case alterations, word and letter substitution, and emoticon handling) and different classifiers are compared in terms of accuracy and performance.

Outside the Twitter-sphere, a different approach is described in (Tchalakova et al., 2011). In this model, a Multi-Domain Sentiment Dataset (Blitzer et al., 2007), containing tagged product reviews from Amazon website, is used as the training data to extract distinctive phrases (i.e., phrases that usually occur in a particular type of document with a pre-assigned polarity) from the processed texts with the ultimate goal of establishing the polarity of the document.

In the field of Business Intelligence (BI), some other works have addressed problems related with the objectives in *Marble* initiative. In (Funk et al., 2008), a supervised machine-learning system is presented to classify texts by ratings. Sentences are tokenized, words are tagged depending on their function and lemmatization is applied. Uppercase and lowercase combinations are also considered while calculating the polarity. Later, (Dey et al., 2010) introduce a mining platform with three main stages: preprocessing, NLP and text mining, also including a dependency extractor to identify relationships between words in a sentence. Some of the techniques employed include phrase grouping, entity extraction, modifiers and synonyms handling, while the polarity calculation depends highly on the word function in the sentence.

Distinctively, Marble pursues a dynamic where time and on-the-fly analysis is considered from the very beginning. The main novelty in Marble and also in the initial experiences introduced in this paper focuses on longitudinal analysis over time, which could be used to detect important shifts on public opinion, and to correlate them with external and internal (to the brand/company/service/product) events on the same time window.

3 DEFINING THE EXPERIMENT

Marble[1] is a Java-based platform with a MongoDB instance (MongoDB Inc. 2014) for tweets storage. The platform performs all the stages of the opinion mining process: *tweets* collection, processing and polarity rating. Finally, Marble also presents the results online. Figure 1 shows the high level architecture of Marble, highlighting its modular nature.

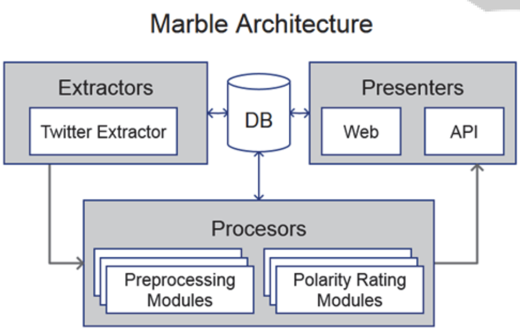

Figure 1: Marble High Level Architecture.

For the particular experiment in this paper, two main topics were selected: the BlackBerry® brand, and the Whatsapp® mobile application. The objective was to evaluate the impact on user's opinion of external and internal events related to both topics. The impact was measured in terms of intensity (number of tweets), polarity change (variation of opinion over time), propagation speed (how fast the event is reflected in Twitter) and duration in time.

[1] Avaliable at: http://iclab.det.uvigo.es/marbleproject.html

3.1 Blackberry Context

On Nov. 04, 2013, Blackberry made an important decision regarding its current directive. Thorsten Heins, who had been the Chief Executive Officer (CEO) of the company since early 2012, was replaced by John Chen. This decision followed a failed takeover offer from an important investment group that affected the company value (Austen and Gelles, 2013).

Another event of minor impact took place on Oct. 29, when the company published the adoption rate for the Blackberry Messenger (BBM) application on Android™ and iPhone® platforms (Bocking, 2013).

An interesting external event occurred on Oct. 31, from one of the main competitors of the Blackberry platform: the presentation of a new Android version (Google Official Blog 2013).

3.2 Whatsapp Context

On Feb. 19, 2014, Facebook announced that it had reached an agreement to acquire Whatsapp for a total amount of approximately $16 billion. The announcement came after months of speculation about which company will acquire it (Facebook 2014).

A few days later, the Whatsapp service was down due to technical issues in their servers. The outage lasted 210 minutes, and also caused problems to Telegram, another mobile messaging service, due to a rush in the service usage as many users installed it to be able to communicate during the outage (Constine, 2014). Although this event was generated inside the company, it can be considered an external event, as it could be caused by external factors (e.g., flood of users, network problems).

4 DATA COLLECTION

Data was extracted by means of GET search/tweets resource of Twitter's Public REST API, which had certain restrictions: extracted data is not exhaustive but a reduced set of the whole twitter universe, as not all tweets are indexed and searchable (Twitter 2013). The extraction module was developed over Twitter4J public library for Java (Twitter4J 2014), and linked to a Mongo DB which stores all the gathered information.

Table 1 shows details about the collected datasets. The collection intervals were selected to cover the events described in the previous sections, and brand names were used as search keywords in Twitter as these brand names are usually associated with themselves. Finally, due to the NLP toolkit, we only extracted tweets written in English.

We made a distinction between original tweets (i.e., tweets authored by the publisher), and retweets. In both cases, the percentage of original tweets is larger than the number for retweets, but in the case of Blackberry (75,6%) it is more acute than in the case of Whatsapp (61%). Also, we extracted the proportion of unique users vs. tweets, and it was quite similar for both cases: Blackberry 55,33% and Whatsapp 54,19%.

Table 1: Datasets Properties.

	Blackberry	Whatsapp
Intervals	2013-10-26 04h 2013-11-06 03h (~ 11 days)	2014-02-15 19h 2014-03-12 20h (~ 25 days)
Keyword	blackberry	whatsapp
Tweets	329.919	2.211.673
Originals	249.532	1.349.162
Retweets	80.387	862.511
Unique Users	182.558	1.198.688
TZ Available	230.723	1.655.247
Unique TZ	141	248

Finally, we were interested in checking the geographic distribution of the dataset using the geolocation fields of each tweet. Unfortunately, this information is not available for an important amount of tweets on each dataset, maybe due to privacy concerns or technical difficulties on a big group of users. As an alternative, the time zone (TZ) used by the user publishing each tweet was available for approximately the 69% of the Blackberry tweets, and for 74% of the Whatsapp ones.

Table 2: Top Time Zones per Dataset.

Blackberry		Whatsapp	
Eastern Time (USA & Canada)	35.590	London	164.832
London	20.925	Eastern Time (USA & Canada)	132.625
Pacific Time (USA & Canada)	18.629	Amsterdam	105.147
Central Time (USA & Canada)	17.127	Pacific Time (USA & Canada)	91.454
Amsterdam	12.065	Central Time (USA & Canada)	82.655
Quito	9.573	Singapore	61.917

Table 2 shows the top time zones of each dataset, that is, time zones with most occurrences. As could be noted, USA and UK are the top contributors on each dataset, which was expected as we selected English as the language of the tweets, but they only contributed 45% of the Blackberry tweets with defined TZ. For the Whatsapp dataset the quantity is significantly lower (31%), and the number of unique TZ (248) indicates a further expanded distribution, and a more global potential impact.

5 DATA PROCESSING

The data processing was divided into two phases: preprocessing and polarity extraction.

First, the tweets were pre-processed. Sentences within tweets were separated applying regular expressions over punctuation marks, all words were changed to lowercase, and invalid characters and words (e.g., punctuation marks, quotations, word with numbers) were removed and substituted with white spaces. After the tweets were separated into individual sentences, each one was tokenized using the Natural Language Toolkit (NLTK).

Next, to assign the polarity of each sentence, a modified bag-of-words approach was used. This model disregarded grammar and word functions inside a sentence, but kept a count on word appearances while preserving the order of them. A polarity valued is assigned to the words and phrases which results of the pre-processing stage. Consequently, the polarity of the whole is assigned as the sum of these individual values (words and phrases). In this process, the dimensional level of the SenticNet 2 corpus (hereafter SenticNet) is used as the source for polarity information of words and phrase. The entire corpus was loaded together with tweets into the same MongoDB instance. SenticNet corpus contained polarity information in three levels: positive, negative or neutral, but also additional concepts like sensitivity, attention, aptitude and pleasantness that could be useful for fine tuning future versions of the classification system (the polarity is in fact derived from these four values).

SenticNet corpus contains not only words but phrases up to four words. Thus, the tokenized words are organized into groups of four. If the group matches a SenticNet group, the polarity of it is the one provided by SenticNet. If it is not the case, the group is progressively reduced by extracting words. The matching process continues to the next not-found word or group of words, and starts again in a four-word group. Finally, the sentence polarity is the sum of the polarity of all the groups found. Figure 2 shows a graphical representation of the matching algorithm, where “n=3” represents the maximum number of words in each group, as indexes start at 0. Once we have the polarity of all the sentences, Tweets having multiple sentences were assigned the average polarity of all the sentences, that is, the sum of the polarities divided by the number of sentences.

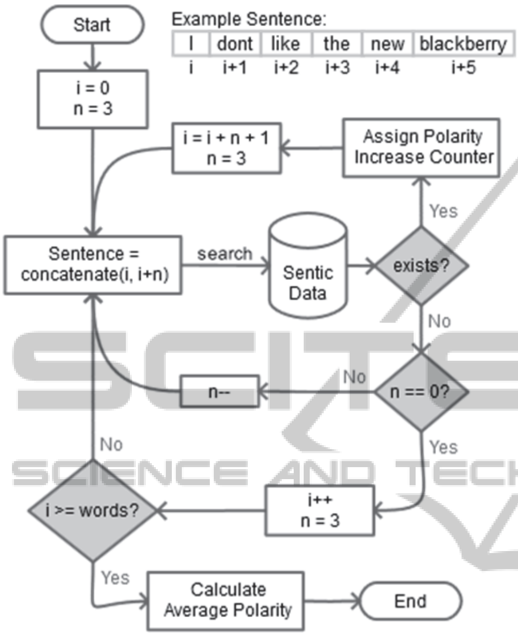

Figure 2: Basic Polarity Calculation Algorithm.

6 FINDINGS

At this point, we were interested in a longitudinal analysis to identify events which impact on users' perception about a brand, that is, analyzing the variances of polarity over time. For that, we defined the “normal polarity” of a brand (keyword in general) by measuring the polarity in a normal period, i.e., a period without relevant events. On this basis, we defined a threshold that allowed identifying a variation of polarity. This threshold of normal polarity was calculated as the average of polarity in the normal period, defined as the sum of every tweet's polarity divided by the number of tweets.

For the Blackberry dataset, we selected two dates to be used as the polarity baseline: Oct. 27 and 28. The average polarity for each one was 0,2511 and 0,2508 respectively, so we selected 0,251 as the threshold value for this dataset.

In the case of Whatsapp, Feb. 16, 17 and 18 were used as baseline dates, each one with average polarity of 0,1406, 0,1506 and 0,1525, respectively. The mean value of these three amounts, 0,1479, was consequently used as the polarity threshold value.

Taking the above thresholds as signs of variation, Tweets were classified as showing a positive (polarity above the threshold), negative (polarity below the threshold) or neutral (polarity between 95% and 105% of the threshold value).

Figure 3 and Figure 4 show the total tweets captured in the defined intervals for each dataset. In the Blackberry case, we observed some unusual tweet counts around the dates selected for the internal events, especially on Nov. 4, where it peaked to 12.000 tweets per hour in our dataset. Its intensity was greater than that of Oct. 29, where its highest value was 6.000 tweets per hour. Also the impact duration on Nov. 4 was a few hours longer than that of Oct. 29. In contrast, no impact was found for the external event on Oct. 31.

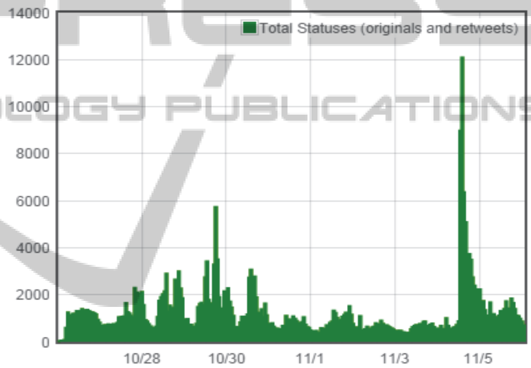

Figure 3: Blackberry Tweets per Interval.

Similarly, Figure 4 shows peaks of traffic on Feb. 19 and 22, the two dates selected for this study. The peak traffic was greater on Feb. 22, but the impact duration was longer on Feb. 19.

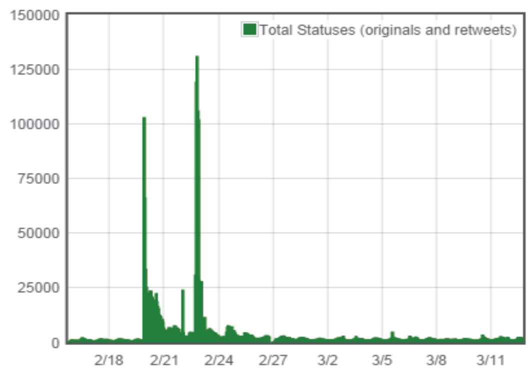

Figure 4: Whastapp Tweets per Interval.

After all the polarities were extracted in each dataset, an hourly ratio was calculated as the sum of all the positive-shift tweets minus the negative-shift ones. Using this ratio, a shift of opinion could be detected using automated mechanisms.

Figure 5 shows the shift ratio for the Blackberry dataset. On Oct. 29 at 13:00 UTC, a positive shift of opinion occurred, but it lasted only 7 hours before returning to the normal situation. On the other hand, on Nov. 4, the opinion shifted toward negative perception, and both the intensity and duration of the opinion shift was greater than the previous event.

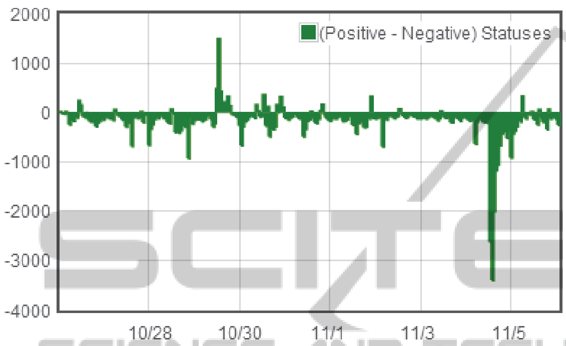

Figure 5: Blackberry Polarity Ratio.

In the case of Whatsapp, Figure 6 shows two main opinion shifts. On Feb. 19, the perception of the Facebook acquisition was positive, and the duration of the positive effect was prolonged for almost two days. At the end of Feb. 22, a big negative opinion shift occurred, just at the same time that the platform was down. As can be seen, the intensity was more than three times the one present on Feb. 19, although the duration of it was shorter.

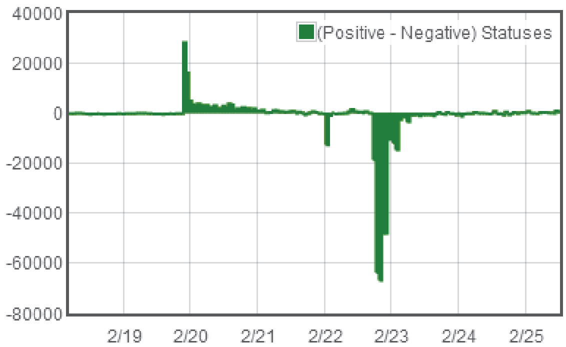

Figure 6: Whastapp Polarity Ratio.

Additionally, we found another shift of opinion not included in the context of the study around midnight of Feb. 22. This was also negative, but it had a very short duration and a lower intensity, compared to the other two. After manually checking the tweets and the news archives, we found out that another outage occurred that day, but affected a lower number of users (Stieber, 2014).

7 DISCUSSION

Information from social networks provides business managers with a valuable resource for making decisions. Precisely, our research approach, Marble Initiative, proposes a methodology that collects relevant data from Twitter (about a single brand or product) to analyze and infer the evolution of users' opinion over time. This information allows business managers to assess the impact on their customers' opinion of internal decision-making and also to detect external events which seems to affect to that opinion.

The data extracted by means of Twitter's Public API, although limited in time and volume, was not irrelevant for the purpose of this study. Moreover, the application of simple preprocessing techniques, SenticNet corpus and a bag-of-words approach provides a fast way to get opinion polarity, which allows a real-time analysis of users' opinion and enables the deployment of an alarms system in the company about perceived image of a product or service.

As the initial launch of the Marble Initiative, the methodology described in this paper provides only a glimpse of all the potential that the system could offer. All the system modules provide plenty of room for improvements, and are being already tested for the next iterations of the platform.

First, the pre-processing techniques could be further expanded. Stemming and lemmatization (Manning et al., 2008) can be used to group similar concepts and avoid getting missing polarities from the SenticNet corpus when the root of the word is in fact present. Synonyms could also be used in cases were the exact word is not found but similar concepts are present. Also the common appearance in Twitter of bad grammar, slangs and text shorthand may be improved by incorporating other NLP techniques.

Second, a disambiguation stage is needed when extracting and processing Tweets. We need to verify that the concept is in fact the one that is being expressed upon and not just being referenced. For example, a user could be talking about how he dislikes something and will review it through his blackberry. Using the described approach, the sentence most probably will have a negative polarity, although the user was referring to something else. Another user could be talking about the blackberry fruit, and his opinion will also be included in the opinion mining results for the brand Blackberry.

The disambiguation of the terms, including the concept verification, is a complex task that requires advanced techniques of natural language processing, but a simple approach, at least for the first example, will be to use the tagging system already incorporated in the NLTK, and identify the keyword function inside the sentence. Techniques like (Michelson and Macskassy 2010) that use Wikipedia as a knowledge base could also be applied.

Finally, for the polarity rating stage, the bag-of-words approach does not handle the effect of modifiers (e.g., not) on the expressed idea, neither the use of complementary sentences that could influence the polarity of the whole Tweet. Both effects need to be included in the rating system, in order to improve its accuracy.

ACKNOWLEDGEMENTS

Work funded by the Spanish Ministry of Economy and Competitiveness under the National Science Program (TIN2010-20797 and TEC2013-47665-C4-3-R); the European Regional Development Fund (ERDF) and the Galician Regional Government under agreement for funding the Atlantic Research Center for Information and Communication Technologies (AtlantTIC); and the Spanish Government and the European Regional Development Fund (ERDF) under project TACTICA. The authors also thank the PhD Programme in Information and Communications Technology from the University of Vigo (Doc_TIC) for supporting travel expenses of this conference.

REFERENCES

Austen, I. & Gelles, D. 2013. *A Takeover Bid for BlackBerry Collapses, and Its Chief Executive Vacates His Post*. Available from: http://dealbook.nytimes.com/2013/11/04/blackberry-abandons-effort-to-sell-itself-c-e-o-to-step-down/ [02 March 2014].

Blitzer, J., Dredze, M. & Pereira, F. 2007, "Biographies, bollywood, boomboxes and blenders: Domain adaptation for sentiment classification", *In ACL*, pp. 187.

Bocking, A. 2013. *BBM – An Incredible First Week on Android and iPhone*. Available from: http://blogs.blackberry.com/2013/10/bbm-first-week/?CPID=SOC_C_WW_TW1383051462 [03 May 2014].

Cambria, E., Havasi, C. & Hussain, A. 2012, "SenticNet 2: A Semantic and Affective Resource for Opinion Mining and Sentiment Analysis.", *FLAIRS Conference*, eds. G.M. Youngblood & P.M. McCarthy, AAAI Press.

Constine, J. 2014. *WhatsApp Is Down, Confirms Server Issues [Update: It's Back After A 210-Minute Outage]*. Available from: http://techcrunch.com/2014/02/22/whatsapp-is-down-facebooks-new-acquisition-confirms/ [03 March 2014].

Dey, L., Haque, S.M. & Raj, N. 2010, "Mining Customer Feedbacks for Actionable Intelligence", *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, pp. 239.

Esuli, A. & Sebastiani, F. 2006, "SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining", *In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC'06)*, pp. 417.

Facebook 2014. *Facebook to Acquire WhatsApp*. Available from: http://newsroom.fb.com/news/2014/02/facebook-to-acquire-whatsapp/ [03 March 2014].

Funk, A., Li, Y., Saggion, H., Bontcheva, K. & Leibold, C. 2008, "Opinion Analysis for Business Intelligence Applications", *Proceedings of the First International Workshop on Ontology-supported Business Intelligence*ACM, New York, NY, USA, pp. 3:1.

Gokulakrishnan, B., Priyanthan, P., Ragavan, T., Prasath, N. & Perera, A. 2012, "Opinion mining and sentiment analysis on a Twitter data stream", *Advances in ICT for Emerging Regions (ICTer), 2012 International Conference on*, Dec, pp. 182.

Google Official Blog 2013. *Android for all and the new Nexus 5*. Available from: http://googleblog.blogspot.ca/2013/10/android-for-all-and-new-nexus-5.html [04 April 2014].

Krikorian, R. 2013*, New Tweets per second record, and how!*. Available from: https://blog.twitter.com/2013/new-tweets-per-second-record-and-how [10 March 2014].

Manning, C.D., Raghavan, P. & Schutze, H. 2008, *Introduction to Information Retrieval,* Cambridge University Press, New York, NY, USA.

Michelson, M. & Macskassy, S.A. 2010, "Discovering Users' Topics of Interest on Twitter: A First Look", *Proceedings of the Fourth Workshop on Analytics for Noisy Unstructured Text DataACM*, New York, NY, USA, pp. 73.

Miller, G.A. 1995, "WordNet: A Lexical Database for English", *Commun.ACM,* vol. 38, no. 11, pp. 39-41.

MongoDB Inc. 2014. *MongoDB*. Available from: http://www.mongodb.org/ [22 January 2014].

NLTK Project 2014. *Natural Language Toolkit*. Available from: http://nltk.org/ [10 January 2014].

Pak, A. & Paroubek, P. 2010, "Twitter as a Corpus for Sentiment Analysis and Opinion Mining", *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, eds. Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, et al, European Language Resources Association (ELRA), Valletta, Malta.

Pang, B. & Lee, L. 2008, "Opinion mining and sentiment analysis", *Foundations and trends in information retrieval*, vol. 2, no. 1-2, pp. 1-135.

Preeti & BrahmaleenKaurSidhu 2013, "Natural Language Processing", *International Journal of Computer Technology and Applications,* vol. 4, no. 5, pp. 751-758.

Stieber, Z. 2014. *WhatsApp Down: WhatsApp Not Working on Friday Night*. Available from: http://www.theepochtimes.com/n3/523179-whatsapp-down-whatsapp-not-working-on-friday-night/ [03 March 2014].

Tchalakova, M., Gerdemann, D. & Meurers, D. 2011, "Automatic Sentiment Classification of Product Reviews Using Maximal Phrases Based Analysis", *Proceedings of the 2Nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis*Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 111.

The Apache Software Foundation 2014. *Apache OpenNLP*. Available from: https://opennlp.apache.org [26 February 2014].

The Stanford Natural Language Processing Group 2014, January. *Standford CoreNLP*. Available from: http://nlp.stanford.edu/software/corenlp.shtml [26 February 2014].

Twitter 2013. *GET search/tweets - Twitter Developers*. Available from: https://dev.twitter.com/docs/api/1.1/get/search/tweets [03 March 2014].

Twitter4J 2014. *Twitter4J*. Available from: http://twitter4j.org/ [03 March 2014].

University of Vigo 2014. *Marble Project*. Available from: http://iclab.det.uvigo.es/marbleproject.html [2014, 23 June 2014].