

Using Social Media for Continuous Monitoring and Mining of Consumer Behaviour

Michail Salampasis¹, Giorgos Paltoglou², Anastasia Giachanou¹

¹Department of Informatics, Alexander Technological Educational Institute of Thessaloniki, Greece, e-mail: cs1msa@it.teithe.gr, giachanou@admin.teithe.gr

²School of Technology, University of Wolverhampton
United Kingdom, e-mail: g.paltoglou@wlv.ac.uk

Abstract. Social communication and microblogging services, together with the ubiquitous online access which keep Internet users constantly connected, provide unprecedented capabilities for the constant connectivity of people and offer a tremendous facility for expressing their opinions, attitudes or reactions about every human activity. In this paper we present an initial investigation of the use of social media for the continuous monitoring and mining of consumer behaviour. We analysed thousands of tweets containing branding comments, sentiments, and opinions about food products. Initial results which are presented in this paper are in line with other results reported in the literature and make us believe that given their dominant use by millions of Internet users, and their distinct characteristics and opportunities, these microblogging services can play a key role in supporting and enhancing important business processes in food industry such as company to customer relationship, brand image building, word-of-mouth branding.

Keywords: Microblogging Social Media, Word of Mouth, Consumer Behaviour, Sentiment Analysis.

1 Introduction

The number of people that use social sites such as Twitter, LinkedIn, Facebook and other microblogging social networking services has grown exponentially the last decade (Huberman et al, 2008). This fact, together with the ubiquitous online access, the increasing adoption of smartphones and the convergence of digital devices, provides unprecedented ground for the constant connectivity of people and offers tremendous capabilities for publicly expressing their opinions, attitudes or reactions about many aspects of everyday human activities. There are so many open questions about the overall impact of these online social services and there is an exponentially increasing interest to study their effects in many sectors and domains from health (Hawn, 2009), to important business processes (Jansen et al, 2009) and education (Junco et al, 2011). In this paper we present an initial investigation of the effects of

Copyright ©by the paper's authors. Copying permitted only for private and academic purposes.

In: M. Salampasis, A. Matopoulos (eds.): Proceedings of the International Conference on Information and Communication Technologies

for Sustainable Agri-production and Environment (HAICTA 2011), Skiathos, 8-11 September, 2011.

social media for the continuous monitoring and mining of consumer behaviour. Using methods from information retrieval, computational linguistics and sentiment analysis, we analyzed roughly 700,000 microblogging messages (so-called tweets) which are mostly dedicated to the discussion of food companies or food products, on a daily basis. We analysed the sentiment of tweets and found some associations with various events related to brand advertisement activities or food policies and crises.

Twitter is a microblogging service allowing users to publish short messages with up to 140 characters, so-called "tweets". Millions of tweets are posted daily by 200 million users, generating 350 million tweets a day and handling over 1.6 billion search queries per day (Twitter engineering blog, July 2011). It is sometimes described as the "SMS of the Internet". These messages usually express personal views of the user, or cover topics such as current political issues, or feelings and attitudes about a new product or a company.

Tweets are visible on a message board of the Tweeter website (<http://twitter.com/>) or through various third-party applications. Users can subscribe to (i.e. "follow") a selection of favourite authors. In such a way one can build a trusted Word of Mouth (WOM) communication channel and can then rely on family, friends and others in their social network to receive opinions about various products. In addition to following a selection of twitter users/authors, and because tweets are also archival and searchable through search engines, users can search for messages containing a specific key word (e.g. a stock symbol, a product which has been recently released, opinions about current political issues). This feature supplies to the opinions which are expressed a certain degree of permanence and increases the effect this electronic WOM, in the form of microblogging, can have in important business processes such brand name development, attention enlargement, reputation management.

The simplicity of publishing short messages on the microblogging services from various communication channels made social websites very popular. A standard tweet is approximately one or two short sentences which make it very easy to create but also to read, making the whole social microblogging service extremely scalable. Internet users prefer microblogs to traditional communication tools such as e-mail. In addition, popular microblogs such as Facebook is often the most visited website on the web, overcoming the popularity of established search engines such as Google (Harvey, 2010). The growing popularity of microblogging services and the content of the published messages led researchers to use them for marketing or social studies (Pak and Paroubek, 2010).

Many of these messages are dedicated to the discussion of food brands and food products. As a result, for customer relationship management or marketing departments of food companies there are tremendous opportunities to obtain useful information they find on social media websites to continuously monitor consumer attitudes, their opinions about brands and reactions about products, and perhaps "understand" their buying behaviour. Twitter based systems have been developed by financial professionals to alert users of sentiment based investment opportunities (Bloomberg, 2010) and by academic researchers to predict break-points in financial time-series (Vincent and Armstrong, 2010). In this paper we aim to present an initial investigation to the issue and to bring attention and discuss a practical application of mining consumer attitude related to food products or brands: what people liked, how they feel and what they would like to see in upcoming products and/or experiences.

For these purposes we would like to examine how these microblogs deliver immediate sentiment which can be monitored and analysed to provide insightful views about products and brands.

In the current research we used Twitter and a system for collecting user generated data (i.e. tweets) for a period of time (approximately one month). We used an algorithm for sentiment analysis proposed by Paltoglou and Thelwall (2011) to automatically analyse the expressed emotion of these microblogs, in terms of a ternary classification scheme, i.e. positive emotion, negative emotion or lack of (objective). The algorithm is based on an unsupervised, lexicon-based approach that identifies sentiment in microblogs and has shown to attain state-of-the-art performance in a variety of social media environments (Paltoglou and Thelwall, 2011).

2 Related Work

Much research has been conducted on sentiment analysis for the domain of blogs and product reviews (e.g. Pang et al., 2002; Blitzer et al. 2007). However, most of this research is based on metadata such as “number of stars”. This is different from extracting opinions from microblogging services. Traditional blogs are very different from microblogs. The main differences between traditional review pages and microblogs are:

- Reviews and text in blogs are longer than the messages that appear on microblogging services.
- Reviews refer to a specific product which is discussed under a topic. On the contrary, the topic of a posted message on microblogs varies.
- Internet users tend to use an informal language that contains non-standard spellings when they post messages on microblogging services (Thelwall and Wilkinson, 2010). This creates a unique and heterogeneous content.

Also, messages on Twitter have characteristics that make them unique:

- *Length*: the maximum length of a tweet is 140 characters. For this reason, research using tweets is different from research based on longer text.
- *Data Collection*: twitter API allows millions of tweeters to be collected for research. That means that the research can be based on real data.
- *Language*: users post in an informal language and it is likely their posts to contain misspellings.
- *Domain*: users post their comments about any topic they want unlike sites that are focused on a specific topic as movies review.

Twitter has attracted much interest the last few years from researchers in the domain of stock prediction and sentiment based investments (e.g. Sprenger & Welp, 2010; Bollen et al, 2011) and several research efforts have been presented in the media (e.g. BBC news, <http://www.bbc.co.uk/news/technology-12976254>) and also resulted in working systems (e.g. <http://tweetrader.net/>) where the real-time sentiment for individual stocks can be accessed. Another utilisation of microblogging

messages is the use of Twitter as a form of electronic word-of-mouth for sharing consumer opinions concerning brands (e.g. Bernard et al, 2009). In this research a large number of tweets are analysed and the researchers conclude that microblogging can be used as an online tool for customer word of mouth communications and could have implications for corporations using microblogging as part of their overall marketing strategy.

The idea behind these research efforts which are based on tracking Twitter is an appealing one: there are millions of people talking about the things they like or don't like or expressing their opinions and attitudes via the Twitter service. Things that they talk about can be books, brands, food, movies, sports or celebrities. If we could analyse the sentiment which is expressed in these microblogging messages, these "sentiment indicators" can be compiled, monitored and analysed by the marketing or other departments within enterprises, or this analysis can be conducted by external service providers specialising in this type of microblogging analysis. As a result very useful information can be timely provided to individuals or organisations such as for example investors, marketing directors, food inspection authorities, which can tremendously help them to improve many of their decisions and processes.

2.1 Sentiment analysis

Profoundly, the problem of sentiment analysis, which aims to determine the attitude of a speaker or a writer with respect to some topic or the overall contextual polarity of a document, becomes very important. Several researchers have focused on the problem of sentiment analysis. Pang and Lee (2008) present an overview of the existing work on sentiment analysis. The authors describe and analyze approaches and techniques that can be used for an opinion-oriented information retrieval system.

Much research has been done for text classification using machine learning. Pang et al. (2002) focused on movie reviews. They tried to solve the problem of sentiment classification using machine learning methods. In particular, they examined the effectiveness of Support Vector Machines (SVMs), Naive Bayes and Maximum Entropy classifiers combined with features such as unigrams and bigrams. They achieved an accuracy of 82,9% using SVM and a unigram model. Turney (2002) tried to solve the problem of sentiment detection by proposing a simple algorithm, the semantic orientation. Pang and Lee (2004) also presented a solution for sentiment detection and analysis. They proposed a hierarchical scheme which first classified the text as opinionated and then as negative or positive. Emoticons such as :-) and :- (have been also used as labels for positive or negative sentiment. Read (2005) used those emoticons to collect training data for the sentiment classification.

2.2 Unsupervised sentiment analysis in social media

Paltoglou and Thelwal (2011) proposed an unsupervised, lexicon-based classifier to address the problem of sentiment analysis on microblogging services. The classifier can estimate the level of emotional valence in text to make a ternary prediction. The proposed algorithm is designed to make opinion detection (i.e. determine whether a text expresses opinion or not) and polarity detection (i.e. whether the opinionated text is positive or negative). The classifier has many linguistically-driven functionalities

that improve its accuracy and suitability for microblogging services. For example, there are negation/capitalization detection, intensifier/diminisher detection and emoticon/exclamation detection. Experiments were conducted for the evaluation of the algorithm. The dataset that was used was extracted from social websites. The results showed that the proposed classifier outperformed machine learning approaches in the task of polarity detection. The lexicon-based approach performed very well in the opinion detection task in the majority of environments. The experimental results showed that the proposed solution is a reliable solution to address the problem of sentiment analysis on social media (Paltoglou and Thelwall, 2011).

The base of their algorithm is the emotional dictionary of the "Linguistic Inquiry and Word Count" (LIWC) software (Pennebaker J. and R., 2001). LIWC contains a broad dictionary list combined with emotional categories for each lemma that were assigned by human. This list was used mainly for two reasons:

1. This dictionary was developed for psychological studies and is extensively used in psychological research.
2. This dictionary is appropriate for informal communication which dominates in social media (Thelwall et al., 2010).

2.3 Systems for Twitter Sentiment Analysis

A number of online sentiment analysis systems have been developed that use Twitter. These systems, like ours which will be presented in the next Section, use the Twitter Search API that can retrieve tweets according to a query term. Most of them are still in the early stages of development and many are focussed on stock market and other sentiment based investments. Some of the most popular applications that can identify sentiment on tweets are presented and discussed below:

- <http://tweetrader.net/>. TweetTrader.net strives to be the most innovative forum for stock microblogs. It is still in the early stages of development.
- <http://www.tweetfeel.com/>. TweetFeel is a free or subscription based service that allows users to track and compare various keywords according to how people use them in Twitter.
- <http://www.socialmention.com/>. Using Social Mention one can receive free daily email alerts (like Google alerts) of a brand, company, marketing campaign, or on a developing news story, a competitor.
- Google alerts (<http://www.google.com/alerts>). Google Alerts are email updates of the latest relevant Google results (web, news, etc.) based on a choice of query or topic.

3. System Description

To experiment with microblogging sentiment analysis, we developed an application that collected data from Twitter services which we aimed to analyze

about their sentiment. The Twitter platform was used to collect the data because it has attracted customer behaviour corporations' interest as it is the largest microblogging service. It is estimated that a million tweets are published every day which reveal customers' opinion and behaviour about companies and brands. Twitter increasing popularity makes it an appropriate source of information to collect data and analyze customer behaviour. That was the reason that Twitter was selected to address the research questions of this paper. The results are expected to be the same if another microblogging service was used as they share a number of similar characteristics.

Twitter provides the Application Program Interface (API) which allows programmatically accessing and retrieving tweets by a query term. The Twitter API takes a set of parameters related to features such as the language, the format of the results, the published date, the number of tweets and a query and returns the tweets that contain the query and meet the parameters. In our system, only tweets written in English were collected as they represent the majority of the published tweets.

An application was developed to collect and store the data. The application queried the Twitter API daily to collect tweets that were posted the day before. Positive and negative tweets were also separately collected with the aid of the Twitter API. For each query, we daily collected the first 1000 tweets, if there are 1000, that were posted the day before. To do so, we used the page parameter as the Twitter API has a limit of 100 tweets for a single request.

We used 35 queries to collect data. The data were stored in a ".json" format to facilitate their analysis. Each json file consists of 100 entries. Each entry refers to a separate tweet. For each tweet the following are stored:

1. Id – id of the tweet
2. Published – date that the tweet was published
3. Link – the link for the tweet
4. Title – title of the tweet
5. Content – content of the tweet
6. Updated – last date the tweet was updated
7. Author – name and uri for the author of the tweet

4. Results & Discussion

The collected data were analyzed for their sentiment. For the sentiment analysis, an unsupervised method was followed (Paltoglou and Thelwal, 2011). Because of time limitations related to the conference submission deadlines, in this paper we present only results from few keywords related to some basic food products such as "milk" (Figure 1), "seafood" (Figure 2) but also brand names ("McDonalds", Figure 3). Figures present the percentage of tweets having one of three sentiments (positive, negative and objective) on the

In Figure 1 we can observe that generally there are 4 times more positive sentiments about seafood in comparison to negative opinions while there are about approximately 50% of them that are characterised as objective (neutral). This simple chart expressing the proportion of negative and positive opinions over a period of time can be a very useful tool, especially if its analysis is made in parallel to

advertisement or other events (e.g. periods that seafood consumption signifies a religious fest).




Figure 1: Seafood

Figure 2 illustrates another interesting example about the power of continuous monitoring of microblogging services such as Twitter. In Figure 2 which depicts positive, neutral and negative opinions about “milk” we can observe a similar pattern of consumer attitude, i.e. there is a relatively stable proportion of positive, negative and neutral opinions based on data retrieved using the word “milk” as the keyword. However, on a specific day (8/4/2011) there is a significant increase in messages expressing negative opinions. A careful examination of this day reveals that was the day of death accidents which reported in China and their cause was related to bad quality or [tainted milk](#). Continuous monitoring can recognize such “abnormal” behaviors immediately and alert food inspection authorities, traceability mechanisms etc.




Figure 2: Milk




Figure 3: McDonalds

Figure 3 represents an example of how Twitter microblogging can play the role of electronic Word of Mouth (WOM). Literature indicates that people appear to trust seemingly disinterested opinions from people outside their immediate social network, such as online reviews (Duana et al., 2008). This form is known as online WOM (OWOM) or electronic WOM (eWOM). This type of eWOM provides consumers tremendous power in influencing brand image and perceptions. For the food industry, in today’s “attention economy” (Davenport & Beck, 2002) where companies compete for the attention of consumers continuously, such sentiment analysis of microblogging data over a period of time can be a tremendous tool to understand how consumers describe things of interest and express attitudes that they are willing to share with others about brands, new food products, concerns about food crises or policies.

Overall, the results show that customers’ sentiment does not change significantly within a small period of time, except when there is a significant reason. This was clearly demonstrated when significant change in sentiment is observed for the query “milk” for which the negative sentiment is increased on 8th of April 2011. However, by tracking, collecting and analysing data over longer period of time, it could reveal much more information about the success or not of important business processes such as brand building.

5. Conclusion

Microblogging has become one of the major types of the communication and this trend is expected to continue. Research indicates Twitter as an important online word-of-mouth mechanism (Jansen et al., 2009). We believe, the large amount of information contained in microblogging web-sites makes them an invaluable source of data for continuous monitoring of consumer behaviour using opinion mining and sentiment analysis techniques.

In our research, we have presented an initial system and method for an automatic collection of microblogging messages from Twitter that are used to understand the opinions and attitudes of consumers towards food products or food related brands. We used an unsupervised, lexicon-based classifier to address the problem of sentiment analysis on microblogging services.

In the future, we plan to expand this work by collecting a larger corpus of Twitter data and analyse how the sentiment analysis of tweets can be used for better understanding of consumer behaviour and how it can be used to improve important business processes in food industry.

References

1. Bernard, J., Zhang, M., Sobel, K. and Chowdury, A. (2009) Twitter power: Tweets as electronic word of mouth. *Journal of the American Society for Information Science and Technology*, vol. 60, no. 11, p. 2169–2188.
2. Blitzer, J., Dredze, M., and Pereira, F. (2007) Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the Association for Computational Linguistics*, Prague, Czech Republic, p. 440-447.
3. Bloomberg, (2010) Hedge Fund Will Track Twitter to Predict Stock Moves, Bloomberg, Online Edition (author Jack Jordan).
4. Bollen, J., Mao, H. and Zeng, X. (2011) Twitter mood predicts the stock market. *Journal of Computational Science*, vol. 2, no. 1, p. 1-8.
5. Davenport, T. H. and Beck, J. C. (2002) The strategy and structure of firms in the attention economy. *Ivey Business Journal*, vol. 66, no. 4, p. 49-54.
6. Duana, W., Gub, B. and Whinston, A. B. (2008) Do online reviews matter? An empirical investigation of panel data. *Decision Support Systems*, vol. 45, no. 3, p. 1007–1016.
7. Go, A., Huang, L. and Bhayani, R. (2009) Twitter sentiment analysis. Final Projects from CS224N for Spring 2008/2009, Stanford Natural Language Processing Group.
8. Go, A., Huang, L. and Bhayani, R. (2009) Twitter sentiment classification using distant supervision. Technical report, Stanford Digital Library Technologies Project.
9. Harvey, M. (2010) Facebook ousts Google in us popularity. *The Sunday Times*, viewed 15th May 2011 <http://technology.timesonline.co.uk/tol/news/tech_and_web/the_web/article70649.ece>.
10. Hawn, C. (2009) Take Two Aspirin And Tweet Me In The Morning: How Twitter, Facebook, And Other Social Media Are Reshaping Health Care. *Health Affairs*, vol. 28, no. 2, p. 361-368.
11. Huberman, B. A., Romero, D. M. and Wu, F. (2008) Social networks that matter: Twitter under the microscope. *First Monday*, vol. 14, p.1–5.
12. Jansen, B., Zhang, M., Sobel, K. and Chowdury, A. (2009) The Commercial Impact of Social Mediating Technologies: Micro-blogging as Online Word-of-Mouth Branding. In *Proceedings of the 27th international Conference Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA. p. 3859–3864.
13. Junco, R., Heiberger, G. and Loken, E. (2011) The effect of Twitter on college student engagement and grades. *Journal of Computer Assisted Learning*, vol. 27, no. 2, p. 119–132.

14. Pak, A. and Paroubek, P. (2010) Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In Proceedings of European Language Resources Association, Valletta, Malta.
15. Paltoglou, G. and Thelwall, M. (2010) Twitter, MySpace, Digg: Unsupervised sentiment analysis in social media. ACM Transactions on Intelligent Systems and Technology, Special Issue on Search and Mining User Generated Contents.
16. Pang, B. and Lee, L. (2008) Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, vol. 2 no.1-2, p. 1–135.
17. 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 of the Association for Computational Linguistics, Barcelona, Spain, p. 271–278.
18. Pang, B., Lee, L., and Vaithyanathan, S. (2002) Thumbs up? sentiment classification using machine learning techniques. In Proceedings of the 40th Conference on Empirical Methods in Natural Language Processing, Philadelphia, USA, p. 79–86.
19. Pennebaker J., Francis, M. and Roger, B. (2001) Linguistic Inquiry and Word Count, LIWC, 2nd ed. Erlbaum Publishers.
20. Read, J. (2005) Using emoticons to reduce dependency in machine learning techniques for sentiment classification. In Proceedings of the Association for Computational Linguistics Student Research Workshop, Morristown, NJ, USA, p. 43.
21. Sprengel, T. and Welpe, I. M. (2010) Tweets and Trades: The Information Content of Stock Microblogs. SSRN working paper, TUM School of Management.
22. Thelwall, M., Buckley, K., Paltoglou, G., and Cai, D. (2010) Sentiment strength detection in short informal text. Journal of the American Society for Information Science and Technology , vol. 21, no. 12, p.2544-2558.
23. Thelwall, M. and Wilkinson, D. (2010) Public dialogs in social network sites: What is their purpose? Journal of the American Society for Information Science and Technology, vol. 61, no. 2, p. 392-404.
24. Turney, P. D. (2002) Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In Proceedings of the 40th annual meeting of the Association for Computational Linguistics, Philadelphia, USA, p. 417-424.
25. Twitter Search Team (2011) The Engineering behind Twitter’s New Search Experience. Twitter Engineering Blog (blog of Twitter Engineering Division), <<http://engineering.twitter.com/2011/05/engineering-behind-twitlers-new-search.html>. Retrieved 2011-06-10>.
26. Vincent, A. and Armstrong, M. (2010) Predicting Break-Points in Trading Strategies with Twitter. Social Science Research Network, Working Paper Series, p. 1-11