Monitoring Opinions in Online Forums – A Case Study from the Sports Industry

Carolin Kaiser and Freimut Bodendorf

***Abstract*—The Web 2.0 provides a worldwide platform for opinion exchange. There is an increasing number of consumers who are expressing their opinions on products. The analysis of online opinions is an important instrument for market research. Due to the abundance of opinions an automated approach is crucial. In this paper, an approach is presented which allows the extraction, aggregation and monitoring of consumer opinions with the aid of text mining. A case study from the sports industry serves as an example for the application and commercial relevance of this approach.**

***Index Terms*—Opinion monitoring, text mining, online forums, online market research.**

I. Motivation

More and more consumers are engaging in online communities where they talk about products. By exchanging opinions they influence their purchasing decisions reciprocally. The large amount of consumer opinions represents a valuable source of information for companies. Online opinions are available at real time and at no cost and can be tracked continuously. Analyzing the development of opinions enables the identification of chances and risks at an early stage and the launching of marketing campaigns in time. An approach for monitoring the development of online opinions is introduced and exemplified by a case study from the sports industry. Consumer opinions on products are first identified by using methods coming from text mining and then aggregated by calculating an opinion index. Afterwards the development of the opinions is examined. This analysis gives insight into the effects of marketing campaigns and external events on opinion formation. Moreover, opinions towards competing products can be compared and the influence on the sales volume can be measured.

II. Related Work

Text mining aims at identifying patterns in texts. It transforms unstructured text into meaningful structured data. Information retrieval, text classification, text clustering, and information extraction are the main subfields of text mining. While past research in these fields focused on mining facts, recent research concentrates on mining opinions. A lot of work deals with the detection of opinions on the Web (e.g. [6], [13], [15]). There are also many papers which track online chatter over a period of time. These studies examine the dynamics of topic propagation throughout weblogs [9], analyze the diffusion of ideas in social networks [10], and enable the tracking of online opinions in weblogs [12]. Several researchers address the task of detecting events in texts. They introduce methods for identifying new events in unstructured texts [11], propose techniques for browsing document collections in order to detect events [16] and present ways of finding novel events in a temporarily ordered stream of news stories [19]. Some papers analyze correlations between online communication and stock movements. These works observe considerable correlations between communication activity in the blogosphere and stock market movement [4], find that postings from stock forums help predicting market volatility [1] and discover that changes in investors' opinions posted to financial forums are closely linked to abnormal returns [18]. There is also some research on predicting the sales volume on the basis of online communication. For instance, [2] and [8] notice that the number of customer reviews is positively associated with Amazon's sales ranking for books. [17] and [13] prove that customer sentiments about movies are good for predicting box office sales. However, an approach which monitors the development of online opinions with regard to marketing campaigns, external events, competing products and sales volume is still missing. Such an approach is crucial for judging the formation of opinions on products.

III. Approach

The approach for monitoring opinions on the Internet comprises three steps. First, opinions posted to an online forum are classified as positive, negative or neutral with the aid of methods coming from text mining. Second, the opinions are aggregated in the form of an index which reflects the overall opinion of the forum towards a product during a certain period of time. Third, the development of the overall opinion is monitored and analyzed. The effects of marketing campaigns and external events are measured, comparisons to competitors are made and the influence on the sales volume is determined. Tracking opinions continuously enables an early detection of changes and risks as well as the initiation of appropriate marketing actions. The approach was applied to adidas' soccer shoe Predator and Nike's soccer shoe Mercurial.

IV. Identification of Opinions

The goal is to identify products and their evaluations in postings. The evaluations show how customers rate products.

Manuscript received March 10, 2012; revised April 16, 2012.

C. Kaiser is with Institute of Information Systems, University of Erlangen-Nuremberg, Germany (e-mail: carolin.kaiser@wiso.uni-erlangen.de)

F. Bodendorf is with Institute of Information Systems, University of Erlangen-Nuremberg, Germany (e-mail: bodendorf@wiso.uni-erlangen.de)

There are three different types of evaluations: positive, negative and neutral. The identification of products and evaluations can be considered as a classification task. Passages in postings are divided into classes according to their linguistic attributes. Since in most cases products are referred to by one or more words, all words of a posting are checked whether they represent a product. Classifications of evaluations are not based on a word level but on a sentence level in order to make allowance for negation or irony. Each sentence is assigned to a class according to its polarity. The classification is based on statistical and linguistic attributes of postings that are gained during the phase of text preprocessing. The text is divided into sentences and words and the grammatical function of the words is identified (e.g. verb).

Support Vector Machines (SVMs) [5] are applied for classifying words and sentences as they allow the processing of many classification attributes and enable the learning of classification rules. As input, the SVMs receive training data consisting of passages with linguistic attributes and class labels. The linguistic attributes depend on the type of passage. Sentences are characterized by the words which they contain. Words are specified by their preceding and following words. On the basis of the training data set, SVMs learn classification rules. The task of the SVMs is to determine the parameters of the classification rules in such a way that the best possible classification is achieved.

TABLE I : CLASSIFICATION RESULTS

	Precision	*Recall*	*F-Measure*
Product	0.9756	0.9438	0.9594
Polarity	0.7326	0.7327	0.7326

For validation purposes, opinions posted to the German online soccer community fussballforum.de have been classified. 407 postings consisting of 2095 sentences were extracted from January to December 2007. All sentences were assigned to one of the three classes "positive", "negative", or "neutral" by a human being. Moreover, all mentioning of the soccer shoes "Predator" and "Mercurial" were annotated. For testing classification, precision, recall and F-measure are measured within a five-fold cross validation. Tab. 1 shows the average classification results which were achieved. While success in detecting products was very satisfying, success in learning opinion polarity was less satisfying. One of the main reasons for this difference is the variation in language. Forum users choose similar words when referring to the shoes in question, but have several different ways in expressing their opinion. Phenomena such as negation and irony make learning opinions even more difficult.

Since products and evaluations are determined independently, they must be linked afterwards. Products get the evaluation of the sentence in which they are contained. If, for example, the mentioning of "Preds" in the sentence "I'm looking forward to playing in my new Preds." is identified as the product and the sentence is classified as positive, then "Predator" is rated as positive.

V. AGGREGATION OF OPINIONS

In order to analyze opinion development, the identified individual opinions must be aggregated. For this purpose, an opinion index *o* is defined which characterizes the overall opinion based on the amount of positive, negative and neutral opinions towards product *i* in period *t*. The formula for calculating the index is as follows:

$$o_t^i = \frac{w*pos_t^i - w*neg_t^i + neu_t^i}{\sum_{i=1}^{n}(w*pos_t^i + w*neg_t^i + neu_t^i)}$$

Legend:
o: opinion index
pos: number of positive opinions
neg: number of negative opinions
neu: number of neutral opinions
w: weight
i: product
t: time
n: number of products

The numerator specifies the overall opinion towards a product over a certain period of time. The amount of negative opinions is subtracted from the amount of positive and neutral opinions. The denominator characterizes the overall opinion of all products in question to ensure comparability between different products. Positive and negative opinions are weighted by the factor *w* to reflect the greater influence of these opinion classes on the overall opinion. The values of the opinion index range from -1 to +1. A high value in opinion index indicates a very positive overall opinion of the product in question in comparison to the other products.

VI. ANALYSIS OF OPINION DEVELOPMENT

A. Analysis of Marketing Campaigns and External Events

The aim of the analysis is to monitor the opinion development of advertised products as well as to explore effects of external events and advertising campaigns on the overall opinion. Opinion development can be illustrated with the aid of graphs. Fig. 1 (left side) shows the opinion development of adidas' soccer shoe "Predator" in the German bulletin board fussballforum.de in the year 2007.

The development of opinions within the German forum is closely linked to the events of the German soccer league. The opinion index is at its peak in May when the German soccer league season ends and the final matches of the German DFB Cup and the European UEFA Champions League take place. The fact that the winning teams of the German DFB Cup and the European UEFA Champions League were both outfitted by adidas, may have accounted for the positive opinion towards adidas' soccer shoe "Predator". During the summer break, in which no soccer matches take place, the opinion index decreases. The beginning of the soccer season in mid-August leads once again to an increase in opinion index in September.

Besides external events, advertisement also influences opinion development. The increase in opinion index in March, May and November may result from adidas' advertising campaigns. In March adidas launched the campaign "Impossible is Nothing 2007". In May the campaign "Predator versus F50" reached its peak and in November the campaign "DFB Kit Launch" was initiated for the European soccer championship.

The positive effect of the advertising campaigns on the increase in opinion index can be evaluated by applying a linear regression analysis which identifies functional dependencies among variables [19]. In this case study, the change in opinion compared to the previous month is used as dependent variable, whereas the external events and the advertisement are used as independent variables. The application of the regression analysis reveals standardized parameter values of 0.6 for soccer events and 0.49 for advertisement (likelihood of error less than 5%). Thus, the positive effects of the soccer events and of adidas' advertising campaigns on the overall opinion in the online forum in question are confirmed. According to R^2, 73% of the variations in change of opinion can be explained by the variations in advertisement and events.

B. Competitor Analysis

The objective of the competitor analysis is to compare the opinion development of one's own product with that of the competing product. This enables an early detection of advantages and disadvantages in competition. In case of disadvantages, measures for improving the product or the advertisement can be taken.

The competitor analysis is illustrated by a graph. By comparing the curves reflecting the opinions on the company's own product and on the competitor's product, it is possible to identify strong and weak points. Fig. 1 (left side) shows the development of the overall opinion on adidas' soccer shoe "Predator" and Nike's competing product "Mercurial" in 2007. When comparing the curves, it becomes apparent that "Predator" is evaluated more positively than "Mercurial" in all months, except for February.

The greatest differences between the opinion indices for "Predator" and "Mercurial" can be found in the months of May, June and September. They result from different advertising strategies connected with different external events. While the opinion about adidas' "Predator" is influenced by national soccer events, the opinion about Nike's "Mercurial" is linked to international soccer events. At the end (May) and the beginning (September) of the German soccer season "Predator's" opinion index increases, whereas "Mercurial's" opinion index does not change much. Instead, "Mercurial's" opinion index reaches its peak in July when the Copa America, the South-American soccer championship, takes place. The victory of the Brazilian team might have contributed to this increase. The Brazilian team was not only outfitted by Nike but also Brazil's famous national player Ronaldo endorsed the soccer shoe "Mercurial".

C. Analysis of Marketing Campaigns and External Events

Sales analysis explores the question whether there are correlations between the opinions towards a product in online forums and the sales volume of this product in reality. If the development of the sales volume follows the development of the opinion with a time lag, companies are able to recognize chances and risks at an early stage and can react accordingly. Sales analysis comprises two steps. In the first step, the developing of opinions and sales volumes are compared with the aid of graphs. In the second step, assumed coherences are evaluated by cross correlation analysis.

Fig. 1 (right side) depicts the development of the opinion and sales volume of the soccer shoe "Predator" within 2007. The curves show a similar time-delayed development. For example, the increase in opinion index in May leads to an increase in sales volume in August. The decrease in opinion index in August causes a decrease in sales volume in November. The opinions in the online forum appear to be correlated with the sales volume with a time lag of three months.

The dependency of one curve on another can be evaluated by cross correlation analysis [2]. The cross correlation coefficient measures the degree to which two time series are correlated at a certain time lag. The cross correlation analysis between opinion and sales volume revealed a significant value of 0.61 for the time lag of three months. Thus, the positive correlation lagging by three months is confirmed. The opinions in the online forums can be used for predicting the sales volume in three months. The same correlation exists between the online opinion and the sales volume of Nike's "Mercurial". The cross correlation coefficient has a significant value of 0.58 at a time lag of three months. It seems to take three months until positive sentiments in online forums lead to the purchasing of the soccer shoes.

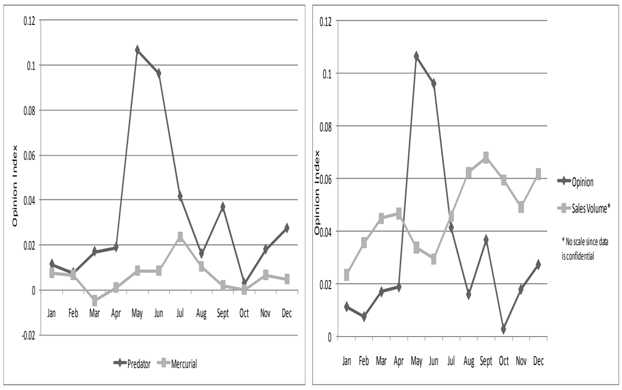

Fig. 1. Competitor analysis (left side) and coherence of opinion and sales volume (right side)

VII. Conclusion

The increasing number of consumer opinions on the Web represents a valuable source of knowledge for companies. The outlined approach allows the monitoring of opinions by employing text mining methods. Opinions are first identified, then aggregated with the aid of an index and finally analyzed with respect to their development. The analysis considers the effects of campaigns and external events, the opinions towards competing products and the influence on the sales volumes. This case study shows the economic relevance of the monitoring approach exemplarily.Future work will extend the basic approach. The aim is to develop an early warning system which recognizes chances and risks automatically by taking the opinions towards a product and its competing product as well as the information about marketing campaigns, external events and sales volumes into account. A warning will be sent to the marketing manager when risks are detected.

REFERENCES

[1] W. Antweiler and M. Z. Frank, "Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards," *The Journal of Finance*, vol. 59, no. 3, 2004.

[2] C. Chatfield, "The analysis of time series: an introduction. 6th edition," Boca Raton: CRC Press.

[3] P. Y. Chen, S. Y. Wu, and J. Yoon, "The impact of online recommendations and consumer feedback on sales," in *Proc. of the 25th Int'l Conf. on Information Systems*, 2004.

[4] M. D. Choudhury, H. Sundaram, A. John, and D. D. Seligmann, "Can Blog Communication Dynamics be Correlated with Stock Market Activity?" in *Proc. of the 19th ACM Conf. on Hypertext and Hypermedia*, Pittsburgh, 2008.

[5] C. Cortes and V. N. Vapnik, "Support Vector Networks. Machine Learning,"vol. 20,no. 3, pp. 273-297, 1995.

[6] K. Dave, S. Lawrence, and D. M. Pennock, "Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Product Reviews," in *Proc. of the 12th Int'l Conf. on World Wide Web*, Budapest, pp. 519-528, 2003.

[7] O. Etzioni, M. Cafarella, D. Downey, A. M. Popescu, T. Shaked, S. Soderland, and D. S. Weld, "A Yates. Unsupervised Named-Entity Extraction from the Web: An Experimental Study," *Artificial Intelligence*, 2005.

[8] D. Gruhl, R. Guha, R. Kumar, J. Novak, and A. Tomkins, "The Predictive Power of Online Chatter," in *Proc. of the 11th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining*, Chicago, 2005.

[9] D. Gruhl D., R. Guha, D. L. Nowell, A. Tomkins, "Information Diffusion Through Blogspace," in *Proc. of the 13th Int'l World Wide Web Conf*, ACM, New York, 2004.

[10] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the Spread of Influence through a Social Network," in *Proc. of the 9th ACM SIGKDD Int'l Conf*, on Knowledge Discovery and Data Mining, Washington, 2005.

[11] G. Kumaran and J. Allan, "Text Classification and Named Entities for New Event Detection," in *Proc. of the 27th Annual International ACM SIGIR Conference*, Sheffield, 2004.

[12] L. W. Ku, Y. T. Liang, and H. H. Chen, "Opinion extraction, summarization and tracking in news and blog corpora," in *Proc. of the AAAI Symposium on Computational Approaches to Analyzing Weblogs*, pp. 100-107, 2005.

[13] B. Liu, M. Hu, J. Cheng, "Opinion Observer: Analyzing and Comparing Opinions on the Web. Proc. of the 14th int'l conf. on World Wide Web," ACM, New York, pp. 342-351,2005.

[14] G. Mishne and N. Glance, "Predicting Movie Sales from Blogger Sentiment," in P*roc. of the AAAI 2006 Spring Symp*, on Computational Approaches to Analyzing Weblogs, 2006.

[15] A. M. Popescu and O. Etzioni, "Extracting Product Features and Opinions from Reviews. Natural Language Processing and Text Mining," Kao A, Poteet S.R. (Eds). London: Springer, pp. 9-28,2007.

[16] D. A. Smith, "Detecting and Browsing Events in Unstructured Text," in *Proc. of the 25th Annual ACM SIGIR Int'l Conf. on Research and Development in Information Retrieval*, Tampere, 2002.

[17] R. Tong, "Detecting and tracking opinions in on-line discussions," in *Proc. of the UCB/SIMS Web Mining Workshop*, 2001.

[18] R. Tumarkin, and R. F. Whitelaw, "News or noise? Internet postings and stock prices," *Financial Analysts Journal*, vol. 57, no. 3, pp. 41–51, 2001.

[19] Y. Yang, T. Pierce, and J. Carbonell, "A study on retrospective and online event detection," in *Proc. of the 21st ACM Int'l Conf,* on Research and Development in Information Retrieval, pp. 28–36,1998.

[20] G. Wang and C. L. Jain, "Regression analysis: modeling &and forecasting," New York, Graceway, 2003.