C. Gurrin et al (Eds.): Advances in Information Retrieval
Proceedings of the 32nd European Conference on Information Retrieval, ECIR 2010
Milton Keynes, UK, pp. 668-669, ©Springer 2010.

Opinion Summarization of Web Comments

Martin Potthast and Steffen Becker

Bauhaus-Universität Weimar, Germany
<first name>.<last name>@uni-weimar.de

Introduction All kinds of Web sites invite visitors to provide feedback on comment boards. Typically, submitted comments are published immediately on the same page, so that new visitors can get an idea of the opinions of previous visitors. Popular multimedia items, such as videos and images, frequently get up to thousands of comments, which is too much to be read in reasonable time. I.e., visitors read, if at all, only the newest comments and hence get an incomplete and possibly misleading picture of the overall opinion. To address this issue we introduce OPINIONCLOUD, a technology to summarize and visualize opinions that are expressed in the form of Web comments.[1]

Related Work Most of the related work pertains to opinion mining in product and movie reviews, where the summarization of reviews has been studied quite intensively [1, 2, 3, 4, 10]. Given a set of reviews on a particular product, the task is to synthesize a summary that contrasts certain product properties a reviewer considers to be positive or negative. In all papers that are referenced here, the generated summaries are lists of ranked sentences extracted from the reviews. Within our approach we focus on words, since extracting sentences is pointless for Web comments: unlike product reviews, Web comments cannot be expected to have a sensible structure or a sufficient writing quality to extract sentences. The difference between reviews and comments becomes apparent if one compares the reviews on products sold at Amazon with the comments on videos published at YouTube. We consider reviews as a special kind of comments, which nonetheless deserve a special treatment. Note further that Web comments in general have been studied far less frequently than reviews [5, 6, 9].

Summarization and Visualization The summarization of a set of comments D divides into an offline step and an online step. Suppose that two dictionaries V^+ and V^- are given, comprising human-annotated terms that are commonly used to express positive or negative opinions [7]. In the offline step we use the well-known sentiment analysis approach described in [8] to extend V^+ and V^- to the application domain. The extension is necessary in order to learn terms that are not covered by the dictionaries. The semantic orientation, SO, of an unknown word w is measured by the degree of its association with known words from V^+ and V^-:

$$\mathrm{SO}(w) = \sum_{w^+ \in V^+} \mathrm{assoc}(w, w^+) - \sum_{w^- \in V^-} \mathrm{assoc}(w, w^-),$$

where $\mathrm{assoc}(w, w')$ maps two words to a real number that indicates their association strength. If $\mathrm{SO}(w)$ is greater than a threshold ε (less than $-\varepsilon$) w is added to V^+ (V^-);

[1] OPINIONCLOUD is available at http://www.webis.de/research/projects/opinioncloud.

otherwise w is considered as neutral. As association measure the point-wise mutual information statistic is applied:

$$\mathrm{PMI}(w, w') = \log_2 \frac{p(w \wedge w')}{p(w) \cdot p(w')},$$

where $p(w \wedge w')$ is the probability of observing w together with w', and $p(w)$ is the a-priori probability of w. In the online step, when a set of comments D is observed, a summary is visualized in the form of a tag cloud which contrasts the positive, neutral, and negative terms found using the sentiment dictionaries. Terms which do not appear in the dictionaries are considered as neutral by default. As is customary for tag clouds, the font size of a term grows proportionally with its frequency in the comments. Moreover, the percentages of positive and negative terms from all non-neutral terms is computed.

Implementation The OPINIONCLOUD is implemented as a browser add-on which, whenever the user views a YouTube video or a Flickr image, downloads the recent comments and summarizes them on-the-fly. The summaries are injected into the Web page. The figures below show examples: the left summary contrasts positive and negative terms on a YouTube video, and the right summary shows the positive, neutral, and negative terms on a Flickr image. For a quick overview it suffices to look at the percentages on top of each cloud which, in this case, indicate that the opinions about the YouTube video are divided with a tendency of dislike, while the Flickr image is clearly appreciated. If a user is interested to know more about what visitors felt when viewing the item, the tag cloud provides the words organized according to their occurrence frequency. By clicking on a word the list of comments containing it is retrieved.

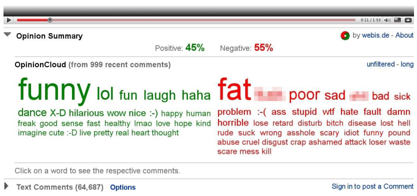

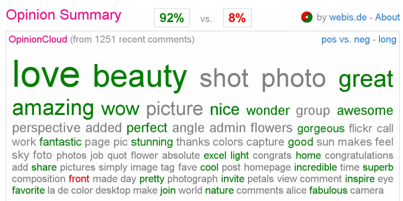

Bibliography

[1] P. Beineke, T. Hastie, C. Manning, and S. Vaithyanathan. An Exploration of Sentiment Summarization. *Proc. of AAAI'03*.
[2] K. Lerman, S. Blair-Goldensohn, and R. McDonald. Sentiment Summarization: Evaluating and Learning User Preferences. *Proc. of EACL'09*.
[3] B. Liu, M. Hu, and J. Cheng. Opinion Observer: Analyzing and Comparing Opinions on the Web. *Proc. of WWW'05*.
[4] Y. Lu, ChengXiang Zhai, and Neel Sundaresan. Rated Aspect Summarization of Short Comments. *Proc. of WWW'09*.
[5] G. Mishne and N. Glance. Leave a Reply: An Analysis of Weblog Comments. *Proc. of WWE'06*.
[6] M. Potthast. Measuring the Descriptiveness of Web Comments. *Proc. of SIGIR'09*.
[7] P.J. Stone. *The General Inquirer: A Computer Approach to Content Analysis*. MIT, 1966.
[8] P.D. Turney and M.L. Littman. Measuring Praise and Criticism: Inference of Semantic Orientation from Association. *ACM Trans. Inf. Syst.*, 21(4):315–346, 2003.
[9] W.G. Yee, A. Yates, S. Liu, and O. Frieder. Are Web User Comments Useful for Search? *Proc. of LSDS-IR'09*.
[10] L. Zhuang, F. Jing, and X.Y. Zhu. Movie Review Mining and Summarization. *Proc. of CIKM'06*.