# Exploring Emotions over time within the Blogosphere


Patrick Hennig
Hasso-Plattner-Institut
University of Potsdam, Germany
patrick.hennig@hpi.de

Philipp Berger
Hasso-Plattner-Institut
University of Potsdam, Germany
philipp.berger@hpi.de

Christoph Meinel
Hasso-Plattner-Institut
University of Potsdam, Germany
office-meinel@hpi.de

Lukas Pirl
Hasso-Plattner-Institut
University of Potsdam, Germany
lukas.pirl@student.hpi.uni-potsdam.de

Lukas Schulze
Hasso-Plattner-Institut
University of Potsdam, Germany
lukas.schulze@student.hpi.uni-potsdam.de



***Abstract*—A lot of research efforts are going on in the area of mining emotions within the world wide web. The BlogIntelligence application is analyzing tons of blog posts and extracts emotions out of this big amount of data. Therefore we thought about how to visualize these emotions in a very meaningful way. While we applied a smart map as a proven technique, we overcame conceptual and technical challenges to provide a feasible utility.**


## I. INTRODUCTION

Nowadays, the society faces a growing importance of information management. More and more eased access to digital media enables people to publish their opinions conveniently and in a high frequency. Namely in customer reviews, blogs or microblogs (such as Diaspora*, Twitter or Facebook) and other platforms of the World Wide Web.

Evidently, there is a strong need also for businesses to acquire publicly available feedback on their own or their competitors products in order to adapt their strategy accordingly. In combination with the evolving impossibility to gather an overview on opinions manually - as it is usual today - reliable and intuitive ways of presenting mined data is of high interest. In addition to the industry, multiple disciplines in science, like political or social sciences, could profit from contemporary analytical tools.

Besides the manifold research on mining sentiments in the vast amounts of user generated content, there seems to be no consensus on how to visualize the results of such analyses. Fortunately, all efforts can now be put in exploring techniques of visualization due to the fact that we experience the availability of fair amounts of documents that are already analyzed for sentiments. Hence, we aim to enable the extraction of yet unidentified information from a visualization.

In particular, a lot of visualization techniques that proved to work well in other domains lack an experimental employment in the field of sentiment analysis. Within something similar like a heat map, the data is represented as colored spots in a matrix. As of today, no major efforts were made to utilize the visual facilities of a heat-like map as a summary of multiple textual documents. Because overlapping spots interfere by design, the map copes the importance of visual aggregation.

Due to the journalistic character of blogs in comparison to microblogs, the retrievable information is expected to be more diverse and semi-structured, ie. because of categories and tags. To feature a reasonable topic-wise presentation, our prototype will rely on data of analyzed blog posts provided by the BlogIntelligence project[1].

To support the use cases mentioned above, the visualized subset of the data must be definable by a search term. Further, we focus on the traceability of sentiments over time, to allow easy recognition of trends and bursts, as explored in [1]. To inform about the context of the visualized sentiments, we will embed the results in related ones, like in results for similar terms. Because users apparently wish natural language in combination with the visualization [2] and want to visit the corresponding source of a specific sentiment, our prototype provides an effortless possibility to enable an interactive exploration. Also the aggregation of resulting data is of relevance, while a histogram aids the user in navigating through the visualized data.

After presenting some selected research in section II that has already been done in this area, we discuss the different frameworks we surveyed to build on in the subsequent section. Another key aspect is described with the analysis of the data to find the best attributes for our dimensions of visualization. This is followed by an explanation of how we developed the emotion map to display the data and to create a more intuitive visualization. Finally we evaluate the user experience and which user interactions are considered useful.

### A. *BlogIntelligence*

With a wide circulation of more than 200 million *weblogs* worldwide, *weblogs* with good reason are one of the most important data streams in the World Wide Web. Therefore, weblogs offer access to latest information discussed in the real world. Since writing posts in weblogs goes along with a high editorial effort, the available information is of major interest. However, for a user it is becoming harder and harder to gain an overview of all discussions in the blogosphere. It is almost

[1]http://blog-intelligence.com/, visited 2014-02-16

impossible for a user to extract information from the web, especially from the blogosphere. Hence, a system that collects information from the blogosphere and presents it to the user in a very meaningful way would be of great use.

Therefore, mining, analyzing, modeling and presenting this enormous amount of data is the overall aim of the project the presented work is integrated in. This enables the user to detect technical trends, political climates or news articles about a specific topic. Most approaches to mining and analyzing such a huge amount of data focus on offline algorithms which use pre-aggregated results. This is in contrast to the continuously growing nature of the World Wide Web. As a result, including the latest data is one of the key aspects of data mining on the web. This is exactly the topic covered by the *BlogIntelligence* project.

The presented work in this paper is integrated into the *BlogIntelligence* project. There are three main steps involved to visualize blogs in the BlogIntelligence project:

*1) Extraction:* In the extraction step the blogs are basically crawled. In order to achieve this a, purpose-built crawler needs to be used as traditional crawlers do not fully meet the particularities of blogs as opposed to conventional websites.

*2) Analysis:* The analysis step prepares the crawled data for visualization. Each blog is analyzed by multiple *Analyzers*, that process its details in certain ways. Among potentially others, there are *data analyzers* that store the meta information about the blogs into the database, *content analyzers* that store information about the content which allow content-related analyses and there are *network analyzers* that store information on the relationships and links between blogs or other communities.

*3) Visualization:* The last step within the BlogIntelligence framework is the visualization of the analyzed information. The Blog IntelliTrends solution is part of this last step as it provides the stored data via an interface and visualizes them in client applications.

## II. Related Work

Neither heat maps themselves, nor the visualization of sentiments are completely novel topics in research and a lot of good publications exist. We limit our résumés to aspects that serve specifically our investigations concerning the visualization of sentiments in a map-alike manner. If a broader overview is of interest, we refer to the excellent and up-to-date section 2 in [3] in all conscience.

It appears that former work had mostly the emphasis on mining sentiments - what is totally legitimate as it is a non-trivial task - and less on visualizing the gathered information. Therefore, the visualization techniques employed until today seem to be limited in their fitness to the verification of the mined sentiments.

The visualization approach of Carter et al. [4] consists of four viewers. In a case study they employ reviews of five products, mined from a major e-commerce platform. The first viewer displays a topic-wise clustered overview of the reviews in the data set. Secondly, a rose plot summarizes the median, the quartile and the relative frequency of occurrence of four pairs of opposite sentiments (ex. Positive-Negative, Pain-Pleasure). The reviews that are visualized in the affect summary (ie. the rose plot) may be selected either topic-wise or product-wise. Thirdly, a histogram for a selected topic helps examining the distribution of the number of reviews per product. And finally, a document viewer can be used at any stage in the process to explore the original reviews. In absence of a user testing, there is no profitable information about the acceptance of the visual metaphors.

A highly beneficial related work was contributed by Pauls and Carenini [2]. Their work focuses ab initio to be generic enough to support diverse information, to feature a programmatic and intelligent selection of relevant information as well as on the development of interactive techniques to provide easy exploration of the data to the user. The core concept of their visualization is a tree map that displays affects using colors. Two zoom steps enable the magnification of a single cell of the coarse-grained tree map to a corresponding fine-grained tiling view. While in the coarse-grained view the color is aggregated from the comprised sentiments, the fine-grained view offers an overview over all non-aggregated sentiments that were mined. Pauls et al. found that - besides superficial critique - a tree map is perceived intuitive and informative by users. Further attention could be payed to the textual summaries and the coordination of text a graphics.

Within the papers examined, the visualization that comes closest to a heat map can be found in a paper about visualization of opinions of hotel reviews [5]. The visualization developed by Nørvåg et al. is a five-dimensional geographical map with a cartographic underlay. Besides the obvious dimensions longitude and latitude; in size, color and opacity varying circles represent the mined opinions. They investigated on how to correlate the subjective dimensions (ie. size, color and opacity of a circle) to the objective data (ex. number of opinions or the numeric representation of an opinion). All historic data get aggregated onto a single map. Due to the fact that a visualization of a time span on such a map is challenge and not within the scope of their work, an additional graph to detect bursts or to reveal long-term changes was implemented.

Another visualization tool of hotel reviews was developed and described by Lin et al. [6]. They built a visual sentiment analysis system - in short VISA - which is based on TIARA, a visually text analytic system. VISA introduces the evaluation and illustration of sentiments based on categories and topics. Hence their visualization is splitted into three views: The trend views, a facet view and a sentiment snippet view. The sentiment snippet view displays the found sentiments with the corresponding text snippet. On the facet view the user gets pie charts with the distribution of measurable values, in this concrete case the customer rating and the type of reviews. The trend view visualizes the found sentiments with their severity and emotion over time. Identified keywords are displayed within the trend, whereas the font size of a keyword depends on the frequency in the corresponding text snippets. The height of each trend stream is mapped to the count of found sentiments and adjusts oneself over time.

## III. Implementation

### A. Frameworks

For implementing our visualization we analysed currently existing frameworks for heat map visualization.

The *High Performance JS heatmap* implementation of Florian Boesch[2] is based on WebGL, an OpenGL pendant for web browsers. In our tests we figured out that WebGL is not fully supported by most of the browsers today and this implementation might not work in many of them.

*JQuery Data Heat Map*[3] is another heat visualization of data. It is developed by Lee Chestnutt and works with the JavaScript-framework jQuery. The data are displayed in a table on which each cell gets an individual color based on the value the cell contains. Because this solution does not support interfering areas, it does not fit our requirements.

The project *heatmap.js* of Patrick Wied[4] uses also jQuery and the native browser implementation of canvas to visualize 2D objects. The canvas standard was introduced with HTML5 and is, same as jQuery, fully supported by the common browsers. The visualization can be done on a map like OpenLayers or on a canvas with an individual background.

On all frameworks we analyzed, *heatmap.js* spotted out to fit best in our scope. It fulfills our criterion to use as much native constructs in our implementation as possible. Furthermore *heatmap.js* and its dependencies are completely open-source which allows us to reuse and modify it in any way.

### B. Data Analysis

We use the data from BlogIntelligence which contains more than 4.6 million different blog posts. On more than 770 000 entries sentiments are identified, whereas the average count of sentiments is 6.53 sentiments per post. The sentiment identification seems to work better for positive ones, because with nearly 4 million positive sentiments. There are four times more positive sentiments in the database than negative ones.

The visualization has three dimensions: x-axis, y-axis and heat. The heat dimension is fixed to sentiments, because we want to visualize the emotion for a specific search term. Each sentiment is referred to a sentence, a paragraph and also to a blog post. Thus it is possible to visualize the sentiments with different granularities like the emotion of a complete blog or the emotion for each post of a blog. Although it is possible to display the emotion in different granularities the best result was achieved by getting the average sentiment for each blog post which keeps the heat visualization clear and understandable for the user.

On the x-axis we decided to use the publishing date of a post, so the user can get an overview how the emotions evolved over time for a specific search criterion. Also the time period can be easily clustered to days, weeks or months to give a better understanding of the emotional development. As already described, the sentiment-analyzed posts constitute a fifth of all crawled posts in the database. Moreover the publishing date of the posts is not distributed equally over time. As a consequence of these, it is quite likely that the amount of sentiment-analyzed posts differs in each time range. Therefore the severity of a visualized sentiment also differs and has a larger impact the less sentiments are available for a specific date [5].

For the third dimension we analyzed attributes which correspond with the search criterion given by the user. By doing this we get an additional connection to the search criterion and the visualized result will be more useful than with any other countable attribute. The database of BlogIntelligence offers an API to get similar terms of a specified term. These similar terms could be used to generate the third dimension and would increase the insight of visualized data [6]. For each search term the interface returns a different count of each similar term. Together with the similar term a correlation rate is returned. It turns out that mostly the first 20 terms correlate in a reasonable way with the given search term. Related to this and for a good visualization we limit the number of displayed terms on the y-axis to 20.

Beside similar terms it could also be efficient to use related tags as third dimension. Single words, manually chosen by the post author, are heavily related to the topic of the post. However the author can decide how to write the tag and because of this the database of BlogIntelligence contains a lot of duplicated tags which differ in capital and small letters as well as in the use of dashes. We analyzed the existing tag data and figured out that mostly the capitalization of letters differed. To solve the difference in most tags, we lower all characters in a tag. To get as many correlated sentiments and blog posts as possible, we decided to use the most frequently used tags for a given search term. First we scan the post content for the search term and collect the top 20 tags which are used with the given search term. Then we collect all sentiments of blog posts which contains the search term in their post content and one of the top 20 tags in their tag attribute.

### C. Emotion Map

As explained in subsection III-A, our prototype is based on the *heatmap.js* project. The subsequent paragraphs explain which challenges we faced during the adaptation of the existing implementation to our needs and how we overcame them. An example emotion map can be found in figure 1.

Initially we determined reasonable parameters and fixed minor defects that originated from *heatmap.js* to get a base to build on efficiently. Firstly, meaningful captions were added to the legend so that an instant orientation concerning the employed colors is possible at any time. Further, we corrected the drawing of spots with a parametrized radius and blur in order to be able to effortlessly alter how spots impact their surroundings. To avoid that spots visually overlay previously added ones, the spots get drawn with a transparency of 50%. By that it is ensured that all spots, regardless of the order of drawing, have a visually equal impact.

To adequately support interference, the aggregation of overlapping spots required a major re-implementation. On a traditional heat map, as implemented in *heatmap.js* , interference solely enabled aggregated values to interfere towards *one*

[2]http://codeflow.org/entries/2013/feb/04/high-performance-js-heatmaps/, visited 2014-02-16

[3]http://www.designchemical.com/lab/jquery/demo/jquery_data_heat_map_demo.htm, visited 2014-02-16

[4]http://www.patrick-wied.at/static/heatmapjs/, visited 2014-02-16

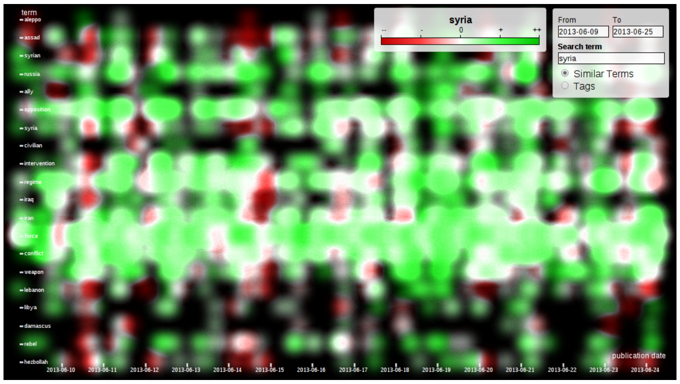


Fig. 1. A emotion map visualization for the search term ”Syria”

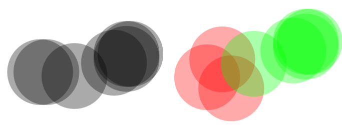

Fig. 2. In contrast to the left method of drawing spots for aggregation, the right method allows to differentiate between positive and negative interference, even within the blended areas.

extrema of the dimension. This is adequate for the domain of thermal energy, because there is no concept of ”negative heat”. For aggregating sentiments, negative sentiments can add up as well as positive ones. As a result the implementation must enable interference towards *both* extrema of the dimension.

The actual interference was computed by drawing shades of uni-colored spots on a virtual canvas. If the sentiment value range is mapped to different shades of that color, multiple interfering weak negative sentiments would add up to positive ones. To eliminate this issue, the superseding implementation computes the correct interference by using shades of two colors. Consequently, the negative sentiment value range is mapped to one color and the positive sentiment value range to the other. The colors must fulfill the condition that they do not have any component of the used color space in common. As an example, in RGB, red and green will never blend indistinguishable. That way it is always possible to determine how much positive and negative sentiments were added. Figure 2 shows schematically the original implementation for computing the interference on the left and the replacement implementation one on the right.

A further issue arises, when two spots stack and their sentiment values are close to one extremum of the value range. The worst case is an aggregated spot with a sentiment value twice the minimum/maximum value. If the emotion map would react to those situations, the value range for drawing the sentiments would need to be extended to include the new extrema. Due to the varied mapping from sentiment values to the corresponding color for the virtual canvas, all already existing spots would need to be redrawn. A questionable effect would be that values close to the original minimum/maximum would appear far less negative/positive. We elaborated not to implement the dynamic expansion of the value range, as single stacked spots can lower the informative value of the emotion map as a whole.

Concerning the colorization, the goal is to provide a meaningful mapping from sentiments to colors on the emotion map using a color gradient. This gradient ranges from dark red, over white, towards dark green. It correlates to the most negative, neutral and most positive sentiment. To retain flexibility, the implementation applies a color mapping (figure 3) to derive the actually visible canvas from the virtual one. It allows a flexible presentation of the aggregated sentiments.

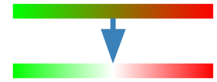

Fig. 3. A color mapping to transform the colors that resulted from aggregation enables the appliance of any color gradient for presentation.

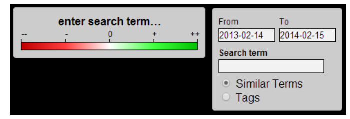


Fig. 4. On the emotion map the user can choose the period, the dimension for the y-axis and he can enter a search term.

Displaying the range of the data is an important information for the user. The original version of *heatmap.js* does not support axis. Therefore we implemented this feature from scratch. The axis work dynamically based on the data used for visualizing the emotion. On the x-axis the range between the minimum date and maximum date of the returned data set is divided equally. All values of the tick labels on the x-axis are converted from milliseconds to days. This ensures a better understanding of the correlation between displayed spots and the position on the x-axis.

On the y-axis we decided to provide two dimensions the user can choose from. The displayed values on this axis are related to the search term and for comprehensibility limited to a maximum of 20 values, which is describe in section III-B. For the tag dimension the tags are ordered from top to down with the most important tag on the top. Contrary to the tags, the similar terms are ordered around the search term which is displayed in the middle of the y-axis. All other terms are ordered around the search term in descending order of their similarity. The highest one is directly on the top of the term, the next one directly under the term, the following again on the top and so on.

### *D. User Experience*

When the user visits the emotion map page for the first time, one of our goals is that the map is self-explaining. The emotion map itself is a central part of the page and has a black background. On the right top corner there are two boxes, one is for the legend, the other one is for the user interactions. As shown in figure 4 the user can choose the period for the visualization. If a date field is focused, a date picker appears to help the user selecting the period. As describe in III-B the user can also switch between similar terms or tags, which defines the dimension criterion of the y-axis. The text field in the middle of the box can be used to enter an arbitrary search term the user wants a visualization for. After entering the search term and pressing enter, the data is requested from the database and will be displayed.

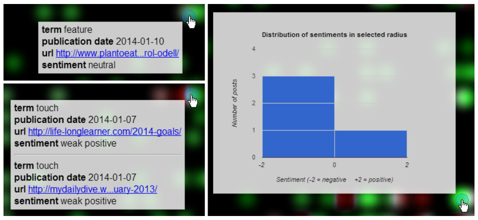


Fig. 5. The text-based tooltip is displayed for less than three overlapping spots. Otherwise a histogram is shown.

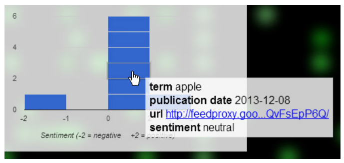


Fig. 6. Hovering one bucket of a histogram displays the text-based tooltip for the bucket.

To improve the users satisfaction the visualized data can be investigated by the user. By hovering the spots additional information for each of them will be displayed in a tooltip. As also mentioned by Carter et al. [4] and Pauls et al. [2] it is helpful to show additional information for a better understanding of the drawn sentiments. Figure 5 shows the behaviour if multiple spots are interfering. Up to three spots are displayed in a text-based tooltip. Beginning with four overlaying spots, the sentiment information is shown as a histogram as visual summary to retain overview [4]. By clicking into the map the currently displayed tooltip will freeze to enable the user to explore the additional shown information. Clicking again on the map unfreezes the tooltip and restarts revealing the spot data on the current mouse position. In figure 6 the mouse was moved over a freezed histogram. When moving the mouse over the blue squares, a normal text-based tooltip is shown. Freezing the second tooltip is also possible and works the same way like freezing a tooltip on the map.

Based on the tooltips and the similar terms as y-dimension the user can investigate how the emotion of a specific term changed over time. Not only the user can detect negative or positive bursts he also can figure out how immense the blogosphere is talking about the specified term and the corresponding topic.

Bursts in the visualization of a product search can reveal at which time a new product version was introduced and whether there was a good resonance of this product the weeks after introduction. Many companies can use this as an indirect feedback channel to improve their products. If the overall perception of their company in the blog community is going bad they can identify the cause of this movement faster and launch countermeasures against it. Using the tag-dimension on the y-axis companies can find out which keywords are mostly used with their product in a specified period. With these information they can create particularly promising advertising campaigns to improve their product success in the market.

In science mostly it is useful to have as much data as possible. The emotion map can shortcut the data analysis and directly provide summarized data. With the additional source of information the researchers can generate more accurate results in a shorter time and they may discover new fields of research based on the results the emotion map provided them.

Even politicians can retrieve additional information for their election campaign from our emotion map. They can keep track of their probability to be voted and compare themselves with their competitors. If another politician becomes very popular, their rivals can determine the date when it started and they can identify and adapt the strategy the popular one uses. If there is a need for a new law they can first check what the population is thinking about the topic to get a better feeling how the law should be formed to be the best for the majority of people in the country.

## IV. Evaluation

To obtain an in-depth understanding of the acceptance and feasibility of our prototype, we employed multiple qualitative user testings.

We surveyed the users expectations regarding the colors that will be employed to visualize sentiments. Our selection of colors for creating the heat map seemed to meet a common cultural understanding of which color is perceived positive or negative. Users independently agreed entirely on using red for negative sentiments and green for positive ones. For the transition between the extrema (neutral), users expressed different colors of preference, such as yellow, white or blue. Nevertheless, the color for the transition was of significantly fewer relevance to the users.

Another aspect we examined in beforehand, is whether users expect to prefer to see aggregated data. Users considered the aggregation of sentiments as a core functionality to survey the affects of a specific topic within the blogosphere.

Further, a majority of the probands expressed the need to browse the original data in order to verify their assumptions.

Summarizing the feedback we received prior biasing the users with the exposition of our prototype, there are strong indicators that their expectations will be met.

The fact that the critique lacked conceptional concerns, confirmed our expectations about the immense potential of this visualization technique. We learned that a few easily modifiable parameters, such as the size of the spots, are not yet optimally chosen and require empirical examination. Some interactions could profit from an enhanced discoverability.

## V. Future Work

In addition to the currently used dimensions, the paper of Burnett et al. [5] discussed adding more information to the emotion map by varying the size of the displayed spots. It can be used to reflect how often a keyword in the context of the values range appears. As well as adapting the size each spot opacity can be changed according to a measurable value. Thus a different meaning can be visualized by having glowing spots. Another approach is to add more colors to the emotion map. In our implementation we use colors only for the extreme values and the neutral value. All spot values are transformed into the corresponding color of the gradient which is generated from our three colors. By adding a specific color for each significant intermediate value the visual feedback can be improved.

Beside the dimensions the displayed values can be clustered. On the x-axis a clustering algorithm can accumulate all values depending on the range of the minimum and maximum date. For instance the clustering could be by day, week or month. However it is important to avoid heavy agglomerations which is more confusing for the user than helpful.

Agglomerations seem to be a current visualization problem if the period is not specified small enough. At times the huge amount of data available on the blogosphere database leads to large overlapping areas on which extreme spot values visually vanish. To solve this problem a threshold could be defined which highlights spots whose emotion value is larger than the difference between the average emotion value of the area plus the threshold. This modification has to be explored carefully since it heavily depends on the quality of the data set.

In the domain of molecular biology heat maps are often used to visualize genes across a number of comparable samples like cells in different states. On these maps the data is mostly arranged to an expressive visualization by hand. After the arrangement the author tries to find meaningful clusters or explanations for the axis. May this approach can be adapted for visualizing emotion maps achieving a comprehensive and meaningful polarized visualization. Such an algorithm seems to be implemented with detailed domain specific knowledge and some kind of machine learning which goes beyond the scope of this discussion. But in general it could be an interesting topic to improve the visualization of data in a heat map.

## VI. Conclusion

At the beginning we introduced some papers of heat map and sentiment visualization. In respect of them we analyzed the given database of BlogIntelligence for suited criteria to use in our map. As also mentioned in the paper of Lin et al. [6], it turned out that the time is a good criterion for the dimension which is mapped onto the x-axis. The dimension of the y-axis can be chosen by the user, but is currently limited to the most used tags and the similar terms for a given search term. While adapting *heatmap.js* as the best matching framework, the scope of the framework moved from heat map to emotion map, because the interference between two extrema was needed for a successful visualization. The user can interact with the emotion map by defining the period of time and choosing the dimension of the y-axis. On the visualized result the user can explore additional information from the spots by hovering them. Overlapping spots are displayed as a histogram if the spots are very dense. In user testings it turned out that we largely meet the users expectations, but extended research efforts in streamlining the user interface could provide a more satisfactory experience. Along with promising research questions that aroused, multifaceted enhancements need further investigations.

## References


[1] G. Mishne and M. de Rijke, "Moodviews: Tools for blog mood analysis." in *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*. AAAI, 2006, pp. 153–154.

[2] G. Carenini, R. T. Ng, and A. Pauls, "Interactive multimedia summaries of evaluative text," in *Proceedings of Intelligent User Interfaces (IUI)*. ACM Press, 2006, pp. 124–131.

[3] A. Das, S. Bandyopadhyay, and B. Gambäck, "Sentiment analysis: What is the end user's requirement?" in *Proceedings of the 2Nd International Conference on Web Intelligence, Mining and Semantics*, ser. WIMS '12. New York, NY, USA: ACM, 2012, pp. 35:1–35:10.

[4] M. L. Gregory, N. Chinchor, P. Whitney, R. Carter, E. Hetzler, and A. Turner, "User-directed sentiment analysis: Visualizing the affective content of documents," in *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, ser. SST '06. Stroudsburg, PA, USA: Association for Computational Linguistics, 2006, pp. 23–30.

[5] E. Bjørkelund, T. H. Burnett, and K. Nørvåg, "A study of opinion mining and visualization of hotel reviews," in *Proceedings of the 14th International Conference on Information Integration and Web-based Applications & Services*, ser. IIWAS '12. New York, NY, USA: ACM, 2012, pp. 229–238.

[6] D. Duan, W. Qian, S. Pan, L. Shi, and C. Lin, "Visa: A visual sentiment analysis system," in *Proceedings of the 5th International Symposium on Visual Information Communication and Interaction*, ser. VINCI '12. New York, NY, USA: ACM, 2012, pp. 22–28.