
Fisheye Word Cloud for Temporal Sentiment Exploration

Ji Wang

Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
wji@cs.vt.edu

Kyle D. Dent

Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
kdent@parc.com

Chris L. North

Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
north@cs.vt.edu

Copyright is held by the author/owner(s).
CHI 2013 Extended Abstracts, April 27–May 2, 2013, Paris, France.
ACM 978-1-4503-1952-2/13/04.

Abstract

This poster abstract presents a new word cloud technique, the Fisheye Word Cloud, for exploring time-series data in a *focused+context* approach to analyzing word data. Our design has two features: cursor-centric layout and word cloud generation on demand. We conducted a validation study to evaluate how our Fisheye Word Cloud influences user performance in comparison tasks of time-series data. Based on task completion time and a TLX-based Likert-style questionnaire, we found the Fisheye Word Cloud has faster *task completion time* and a better *user satisfaction level* than the alternative we reviewed.

Author Keywords

word cloud, user study, sentiment analysis, temporal twitter data analysis

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous.

General Terms

Human Factors; Design; Measurement.

Introduction

As social media has become one of the dominant methods for individuals to communicate and express themselves,

companies are anxious to be involved in the conversation. People are talking publicly about products, services, and the brands they interact with. In the past companies controlled their own message. Now consumers are exposed to a lot of opinions, and they're probably more influenced by their peers than by corporate marketing campaigns. Companies are interested in trends, especially sentiment changes over time. Our business intelligence project is helping companies to understand this new data and incorporate it into their existing market analyses.

For this project, we wanted to use a word cloud visualization not only to present words but also to understand the sentiments associated with them over time. As we saw it, we faced three main challenges in presenting data for this kind of comparison:

- How can we present temporal changes in a word cloud? Can we also show trend information over time?
- How can we let users interact and explore the word cloud contents in time-series data, and let users get a sense of the trending of the most frequent words?
- How can we deal with the scalability problem inherent in word cloud implementations for time-series data?

To address the challenges, we created an interactive word cloud that we call the Fisheye Word Cloud. Words in the cloud that are closer to the focal point are larger and more easily distinguishable to users. The level of display detail decreases as words move further away from the center creating an effect that provides insights into the words people are using most frequently without losing the overall context for understanding the bigger picture. Our word cloud permits sentiment exploration and provides two main strengths:

- Scalability, even for viewing data over various time periods and at different granularities of time.
- User-controlled focus, centering the features analysts are most interested in and tracking changes over time.

Certain types of word clouds, such as the Parallel Tag Cloud, can present sentiment changes over time efficiently, but as more time is added they need larger and larger space. Our goal was to use fixed space to display time-series data in a straightforward way, and at the same time embed the sentiment for that time thereby allowing users to focus on a particular time.

The contributions of this poster abstract are (1) to explain the design of the new word cloud for time series data exploration and (2) to present a user study showing evidence that the fisheye word cloud (a) provides a high-level of data scalability; (b) requires lower *physical demand* for the time-series data exploration task; and (c) elicits a better *user satisfaction level* than alternatives we reviewed.

Related Works

Word clouds have several beneficial qualities. They are applicable to a wide range of applications because they can be compiled based on word frequency from any kind of text source. Another beneficial aspect is that word clouds enhance a viewer's insight with an aesthetically pleasing presentation where different font sizes are coordinated to create an attractive visualization compared to blocks of text or lists of words [10]. There are several examples of projects that have made good use of word clouds for presenting information: Review Spotlight [12], Wordle [2] [7] [10], TIARA [9] and for example. At the same time, the random layout of word clouds is a

drawback. Since the layout is random, it doesn't provide context, and if a word cloud is re-generated, the layout is not the same as it previously was. The lack of a consist layout requires a greater mental demand when locating specific words. Also, randomly structured clouds can encode only a single detail, i.e. frequency. Users can determine relative word frequency but without any contextual information about the content the word cloud represents.

used to visualize time-series trends [8]. The SparkCloud embeds trend information, as opposed to frequency information, to provide users a general sense of each word's temporal trend over time. Koh et al added interactions to Wordle visualizations [10] and created ManiWordle [7]. ManiWordle allows words to be moved around, and as words are moved, other words reorganize themselves automatically.

Design

Based on our survey and discussion of currently existing word cloud visualization methods for time series data exploration, we think word cloud visualizations are a good method for comparison tasks. They provide an overview in a fixed space and allow interaction plus the possibility of embedding at least one more dimension of data beyond frequency. But, simple word cloud approaches are not sufficient to support our task for comparing words expressing sentiment over time. To achieve our goals we set two design objectives: (1) Embedding sentiment analysis results (one more dimension of data) into a word cloud, and (2) Providing interactive navigation to support a *context+focus* approach.

Data Collection and Natural Language Processing
To generate our fisheye word cloud, first we collect twitter data from a real-time stream. After some linguistic pre-processing, we extract features from the tweets, which are stored in the system. For each tweet from the pool of collected data, we process the features using a previously trained model for both topic and sentiment detection producing a negative or positive (or neutral) classification for each one. We then compile lists of words and mark them as positive or negative depending on which polarity is assigned to tweets where they appear most often. Using this data we can create the interactive word cloud for

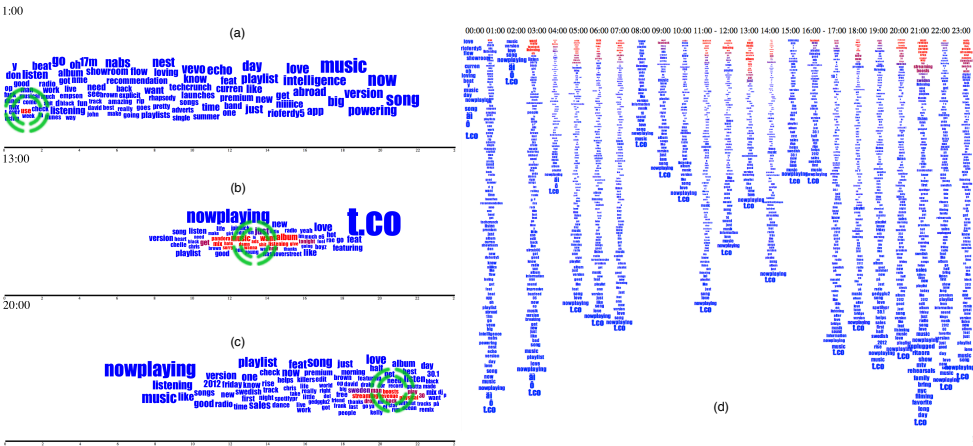


Figure 1: (a) (b) (c) Fisheye Word Cloud and (d) Parallel Tag Cloud. The top-left corners of (a) (b) (c) shows the current cursor's time points in the Fisheye Word Cloud; The bottom axes in (a) (b) (c) indicate time. The green target shows the position of the mouse cursor. The timeline is at the top of (d) in the Parallel Tag Cloud.

There have been some enhancements to word clouds for temporal data visualizations. Collins and Viegas presented the Parallel Tag Cloud to aid users in analyzing and navigating text sources over different time points or different dimensions [1]. The idea of the Parallel Tag Cloud is driven by the Parallel Coordination visualization approach. Lee et al presented the SparkCloud, which is

time-series twitter data exploration, as shown in Figure 1 (a)(b)(c).

Interactive Fisheye Word Cloud Design

Following the linguistic pre-processing and sentiment analysis, we generate a word cloud similar to that of Wordle [4]. The difference is that the initial position is not random, but follows an Archimedean spiral curve. The center part of the mouse cursor is located among the most negative words and the outside contains the most positive ones. The rendering process is handled by 'twisting' this structure of words in a manner akin to twisting a Swiss roll.

Implementation

We implemented the fisheye word cloud in JavaScript and Java. For the natural language processing part, we use tools from the Stanford Parser and OpenNLP projects. For the word cloud visualization, we use D3.js. [3] [11]. We generated two types of word clouds for our user study: our own fisheye and a parallel word cloud for comparison.

User Study

We conducted a follow-up validation study to assess the effectiveness of the fisheye word cloud for business intelligence and sentiment analysis compared to the Parallel Tag Cloud. In our study, the evaluation metrics can be divided into two aspects: *task completion time* and *user satisfaction level*.

For the study, there were 14 participants (5 female). Their average age was 36. The group was comprised of students, researchers, software engineers, business administrators, and account managers. All participants were volunteers and came from within our organization and the local area. Participants used a 3.0 GHz Dell Desktop with 4GB RAM and a Samsung 27-inch LCD

display running at 1920×1200 pixel resolution for visualization tasks which were run on the Google Chrome web browser. Task completion time was measured by stop watch. After performing the assigned tasks, participants completed a TLX-based Likert-style questionnaire and survey on a Lenovo Thinkpad T400s laptop. The user study data was one day (24 hours) of Twitter data. For the data content we extracted the top 100 words from a Twitter collection created by pulling data from our real-time stream using the keywords "Pandora" and "Spotify". The time unit or data granularity we used was one hour.

We implemented a parallel tag cloud as an alternative approach for comparison in our study based on [1] in JavaScript with the D3.js library viewed with a browser, as shown in Figure 1 (d). We encoded each word's sentiment result from top to bottom with the top part containing the most negative words in each time frame, and the bottom part the most positive words in each time frame. The fisheye word cloud was generated in the same way. To adjust the font size users can press ctrl +/- to zoom in/out for comfortable reading. At the same time, there is no interaction embedded in Parallel Tag Cloud and users need to slide left/right and scroll up/down for navigation if the content extends beyond the size of screen.

The study task was designed from a real-world sentiment analysis task: What is the public's opinion towards a specific brand or product. In order to let participants complete this task more easily, we divided the task into three sub-tasks. Task 1 was the overview task. Tasks 2 and 3 were specific data tasks similar to [6] and [8], which used similar types of tasks for their temporal visualization evaluation task design.

- **Task 1 (T1):** What conclusions can you draw about

people's attitudes towards streaming music?
(Overview task)

- **Task 2 (T2):** Describe any aspects of streaming music that people find particularly annoying?
(Specific data task)
- **Task 3 (T3):** Describe any aspects of streaming music that people are particularly pleased with?
(Specific data task)




Figure 2: Task Completion Time of 3 Tasks (in seconds)

In our study, each participant used both of the two different visualization approaches. The ordering of each type of word cloud was counter balanced. At the end of the user study, each participant described the strategies they used for these two different approaches and provided general feedback. There was a follow-up questionnaire with six TLX-based, Likert-style questions [5] (1-10 Likert scale, where lower is better).

Results

We collected data from 14 participants and applied *paired t-test* for the *task completion time* analysis using the square root of the *task completion time* (in seconds) for each task for each participant—a widely-used transformation to correct for the non-normal distribution of this type of time performance data. For the *user satisfaction level* analysis, we calculated the pairwise Wilcoxon test for each TLX-based, Likert-style question for the two word cloud visualization types.




Figure 3: Post-trial survey results from TLX-based, Likert-style questions (1-10 Likert scale, where lower is better)

Task Completion Time: With paired *t*-tests, we found there is a significant difference in the *task completion time* in Task 2 (specific data task: find negative words) for the fisheye word cloud and the parallel tag cloud conditions; $t(13)=2.619$, $p=0.021$. However, there is no significant difference on the *task completion time* in Task 1 (overview task) and Task 3 (specific data task: find

positive words) for both word cloud visualization approaches. The results of the *task completion times* of the three tasks in our study are shown in Figure 2.

User Satisfaction Level: With the pairwise Wilcoxon tests for each TLX-based, Likert-style question for the two word cloud visualizations, we found the fisheye word cloud has a significantly lower *physical demand* than the parallel tag cloud ($p=0.009$). There is no significant difference for the other factors. The results of the six TLX-based, Likert-style questions of in our study are shown in Figure 3.

Discussion

The significant difference in the *task completion time* is due to the fact that negative words follow the cursor as it moves (the fisheye) making it easy for users to pay more attention to the area of the cursor as they explore the data through time. With the parallel tag cloud, users spent more time scanning left and right to get all of the necessary content. The non-significant differences in Task 1 and Task 3 show that users spent time locating words outside the focal point located by the cursor in the fisheye word cloud similar to the time spent scanning left and right in the parallel tag cloud.

The *user satisfaction level* results indicate that the fisheye word cloud compared to the parallel tag cloud can significantly reduce the overhead of physical navigation during the temporal data exploration task.

For small amounts of data or shorter time spans, neither the fisheye nor the parallel word cloud obviously performs better than the other. However, the fisheye word cloud is clearly more suited for larger scale time-series data. The fisheye word cloud can easily visualize 365 days of temporal tag cloud data, while the parallel cloud becomes

unreasonably wide to handle the same data. The long column length is another drawback which leads to a higher physical demand for parallel clouds. The fisheye word cloud provides a more compact word cloud view for each time frame and reduces the overhead of scrolling up and down for data navigation.

The scalability feature of the fisheye word cloud is self-evident. In our study for same dataset, the fisheye word cloud used only 1800×300 pixels for visualization while the parallel tag cloud used 1800×1000 pixels to display.

Conclusion and Future Work

In this poster abstract, we have described the design for a new interactive, cursor-centric word cloud layout. The fisheye word cloud provides a highly scalable presentation for temporal data exploration. The results of the user study we conducted demonstrate that compared to the parallel tag cloud, the fisheye word cloud allows a significantly faster *task completion time* for focus words and significantly lower *physical demand* during navigation over different time periods.

Future plans for our system include making use of this layout and interaction approach to encode more information into word cloud visualizations.

References

- [1] Collins, C., Viégas, F. B., and Wattenberg, M. Parallel Tag Clouds to explore and analyze faceted text corpora. In *2009 IEEE Symposium on Visual Analytics Science and Technology*, IEEE (2009), 91–98.
- [2] Cui, W., Wu, Y., Liu, S., Wei, F., Zhou, M., and Qu, H. Context-Preserving, Dynamic Word Cloud Visualization. *IEEE Computer Graphics and Applications* 30, 6 (2010), 42–53.
- [3] D3.js. <http://d3js.org/>.
- [4] Feinberg, J. Wordle. In *Beautiful Visualization*. 2009, ch. 3.
- [5] Hart, S., and Staveland, L. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human mental workload*, P. A. Hancock and N. Meshkati, Ed. North Holland Press, Amsterdam, 1988.
- [6] Javed, W., McDonnell, B., and Elmqvist, N. Graphical perception of multiple time series. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 927–34.
- [7] Koh, K., Lee, B., Kim, B., and Seo, J. ManiWordle: Providing Flexible Control over Wordle. *Visualization and Computer Graphics, IEEE Transactions on* 16, 6 (Nov. 2010), 1190–1197.
- [8] Lee, B., Riche, N. H., Karlson, A. K., and Carpendale, S. SparkClouds: visualizing trends in tag clouds. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 1182–9.
- [9] Liu, S., Zhou, M. X., Pan, S., Song, Y., Qian, W., Cai, W., and Lian, X. TIARA : Interactive , Topic-Based Visual Text Summarization. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3, 2 (2012).
- [10] Viégas, F. B., Wattenberg, M., and Feinberg, J. Participatory visualization with Wordle. *IEEE transactions on visualization and computer graphics* 15, 6 (2009), 1137–44.
- [11] Word Cloud in D3.js. <http://www.jasondavies.com/wordcloud/>.
- [12] Yatani, K., Novati, M., Trusty, A., and Truong, K. N. Review spotlight: a user interface for summarizing user-generated reviews using adjective-noun word pairs. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press (New York, New York, USA, 2011), 1541.