Interactive Folksonomic Analytics with the *Tag River* Visualization

Angus Graeme Forbes, Basak Alper, Tobias Höllerer, and George Legrady

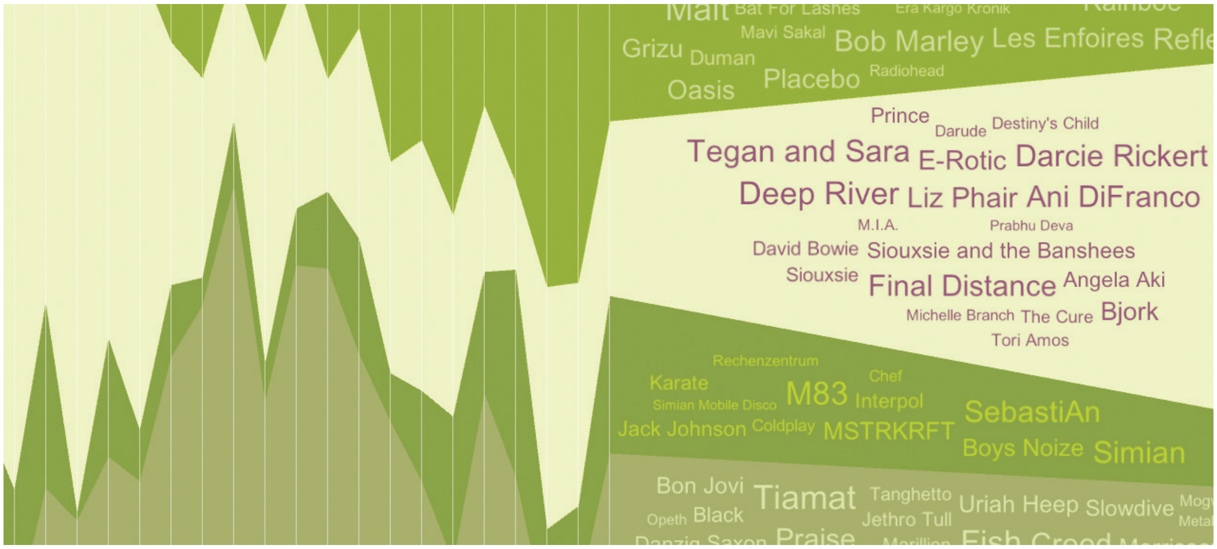

Fig. 1. The Tag River visualization depicting data regarding user's listening habits from the Last.fm social networking site.

Abstract—*Tag River* is a novel visualization that presents a detailed comparative overview between user content for a particular span of time. Simultaneously it provides a trend summarization of earlier or later time spans. The summarization is displayed using vertically-adjacent polygonal regions in which the area represents some facet of quantitative information. A series of animated tag clouds are used to describe more detailed content for each user, changing over time to provide an indication of the coherence of context between time segments. The concurrent representation of both multivariate and temporal data can be cycled though programmatically or navigated interactively, allowing the user to explore time spans via filtering or zooming. Changing the view to a new time span instantly updates the tag clouds. We use color and size to represent information associated with the tags, and these aspects are updated to reflect changes in information when a new time span is selected. To facilitate these updates, we introduce a 2D packing algorithm which satisfies specified aesthetic criteria and runs at real-time frame rates. This paper describes the visualization technique in detail and presents example visualizations using datasets from social media sites.

Index Terms—Text visualization, text analytics, social media, social network data, information visualization.

1 Introduction

Tag clouds have become highly popular over the last few years, outgrowing their initial use as a navigation tool and instead appearing more commonly to indicate content aggregation, high-level overviews, and ad hoc topic summarization via word frequency visualization. Our novel interactive visualization system, *Tag River*, uses tag clouds to provide further detail about a specific time span within a larger temporal dataset. That is, rather than being used to provide a single high-level overview, we appropriate tag clouds to present a comparative sampling of information at a particular segment in time between multiple categories. Moreover, we animate the appearance and disappearance of particular tags within the clouds to indicate continuity (or lack of continuity) over time. Our goal in creating this visualization is to present a methodology with which to compare semantic aspects of social data between categories while maintaining a more definitive quantitative context from which to reason about the social data. Quite literally, we place the semantic, textual and/or cultural information (tags and images) within the boundaries of a polygonal region whose height directly maps to a quantitative datum. Further contextualization is provided by adjacent polygons in the form of stacked bar charts, indicating information at an earlier or future time. In section 4, we present example interactive visualizations based on social media data for exploring how tags associated with particular users change over time. Through interaction with Tag River visualization, users are more easily able to make comparative inferences regarding these users by examining folksonomic data both at a particular time span and as the data changes over time.

- *Angus Graeme Forbes is with the Media Arts & Technology Department at the University of California, Santa Barbara, E-mail: angus.forbes@mat.ucsb.edu.*
- *Basak Alper is with the Media Arts & Technology Department at the University of California, Santa Barbara, E-mail: basakalper@gmail.com.*
- *Tobias Höllerer is with the Department of Computer Science at the University of California, Santa Barbara, E-mail: holl@cs.ucsb.edu.*
- *George Legrady is with the Media Arts & Technology Department at the University of California, Santa Barbara, E-mail: legrady@arts.ucsb.edu.*

Manuscript received 31 March 2011; accepted 1 August 2011; posted online 23 October 2011; mailed on 14 October 2011.
For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

2 Background

Tag River is, in a sense, a hybrid between two popular visualization techniques. Similar to the *ThemeRiver* visualization, introduced by [10], we employ continuously connected stacked regions to visualize temporal information. In fact, Tag River could be thought of as a space-fitting specialization of the ThemeRiver visualization for temporal tag clouds or other folksonomic datasets.

The efficacy of various approaches to tag cloud visualization has been examined in a number of studies. For example, [11] and [9] utilize graph-drawing algorithms to cluster associated tags. Other approaches, such as [2] suggest using circular tag layouts to imply relationships inherent in social networking navigation systems, spatially centralizing significant tags to indicate the defining semantic concerns of a social group. A more recent work explores other ways to preserve clusters of semantic similarities within a word cloud via a seam-carving technique. With this technique, clusters that have a strong relationship can be visualized via spatial proximity or by using bubble-sets to demarcate more significantly related words [22]. Tag River utilizes a layout based on an aesthetic constraint that requires more significant (and hence larger) tags be placed in central positions given an arbitrary polygonal region. Instead of focusing solely on semantic relations within the tag cloud, we instead emphasis context preservation of tag cloud membership over time.

The visualization of changes in semantic data over time via tag clouds has also been proposed by various researchers. For instance, [4] introduces an algorithm that samples salient tags within arbitrary time frames, which are then enhanced with animations that show the movement of tags as they change over time. A paper by [17] describes the evolution of individual tags via a visualization of line graphs. Work by [3] aims to preserve the semantic coherence and spatial stability of tag clouds over time and also includes an alternative trend chart. The recent SparkClouds technique aims to combine the aesthetic approachability of tag clouds with trend information by adjoining each tag with a spark-line to indicate historical information [14]. Our use of semantic tags embedded within data streams is most similar to the TIARA visualization, introduced in [15] and further extended in [21] and [18]. In that system, subsets of text (gathered from the results of a topic modeling algorithm) are placed throughout the visualization to provide an initial summarization of the data streams at different points in time. Tag River approaches the display of semantic information in a somewhat different manner. Tag River constrains semantic information to a single time span at any given time, but allows the user to quickly change the display to view different time spans and also maximizes the amount of information inside the current time span by expanding the current time span to take up more screen space. Animation and shifts in text color are used to interpolate between two tag clouds at two different time spans, providing a rich indication of temporal changes in semantic data.

3 Overview of Visualization System

The Tag River visualization method allows users to examine trends in temporal-categorical data sets. It consists of the following interconnected components: the overview visualization of trends over time; a more detailed tag cloud comparison within a particular time span; the visualization and animation of context preservation between time spans; and real-time interactivity to explore the information at selected time spans.

For a given categorical data set over time, the Tag River visualization divides a rectangular display area into polygonal layers, stacked vertically atop each other. Each of the stacked layers represents a particular data category. The layers are color-coded in order to provide clear visual differentiation between them. The space is further divided into a series of time steps along the horizontal axis. The screen is thus split into a set of polygons, each of which represent a particular category during a particular time span. The boundaries of these polygons are determined by the relative magnitude of a particular datum associated with each category at a particular time. For the example projects (described in section 4), we use the number of semantic tags associated with users of a social media service as our boundary-defining datum. Relative magnitude for each data category is calculated at each time step by dividing the number of tags associated with this category into the overall number of tags for that time step, yielding a normalized measure for the magnitude. The only restriction on the number of data categories is the available screen space, although including too many of these may make it more difficult for a viewer to make meaningful comparisons between categories.

This relative measure is then used to determine the vertical length of the right and left sides of the rectangular region for the previous and the current time step, respectively. Connecting these points with straight lines creates the horizontally-stacked quadrilaterals that fit in the rectangular slice of the screen representing a time span. The position and size of the quadrilaterals conveys a high-level comparison of how the relative magnitude of the different categories change over time. Thus, users, at-a-glance, have an indication of the general trends within the temporal-categorical data set.

Each quadrilateral defines a bounded region in which we display tag clouds containing associated semantic information for that category. These tag clouds provide more detailed presentation of semantic information specific to a particular category during a particular time span. Each tag is scaled in size to indicate the frequency of occurrences of that tag within the current time span. Color transparency is used to indicate the temporal coherence of the tags: tags that occur for the first time within the temporal data set are made semi-transparent, but become increasingly more opaque as they continue to occur successively in the following time steps. The tags are then positioned using a bin-packing algorithm which has been customized to take layout properties into consideration (described below). Since there may be a large number of tags associated with a particular category, it may not be possible to display each tag at a legible size. For this reason we set a lower bound on the size of the tag, which excludes the tags that appear with only a low frequency.

Figure 1 shows a detail of an example Tag River visualization where the right-most time span has been selected as the current time span. It contains four differently colored horizontally-stacked quadrilaterals, each representing data for a different data category during that time span. For instance, at the beginning of the time span the second category from the top of the screen (the strip color-coded white) is associated with a number of tags that is approximately equal to the number of tags associated with each of the other categories. And therefore the left height of the white quadrilateral is approximately one quarter the size of the height of the entire screen. However, by the end of the time span, this category has more tags than all the others combined, and so the right height of the quadrilateral is increased accordingly.

When a new time span is selected (either via interaction or programmatically), the visual elements representing past time steps are modified. The tags shrink and disappear and the polygonal regions are scaled in width, contracting into stacked bar charts that represent the normalized volume of tag information. As in the expanded, detailed view containing the tag cloud, each contracted slice of time contains a set of color-coded vertically-adjacent quadrilaterals filling up a rectangular slice of the screen. The left side again represents the beginning of a time period, and the right side represents the end of that time period. However, the width of these vertically-adjacent quadrilaterals is much thinner than when in the expanded state. Both Figure 1 and Figure 3 show summarization rectangles at time spans surrounding the current time span.

3.1 Context Preservation through Animation and Interaction

In addition to providing trend summarization (via the surrounding time spans), we use the movement of particular tags within tags clouds to indicate the coherence of particular features between time steps. Tags that are present in both the previous time step and the current time step are repositioned and resized through an animation, making it obvious how the prevalence of the tag changes. On the other hand, if a tag from the previous time step is no longer included in the current time step, or if a new tag is now included that wasn't present before, then we animate the appearance or disappearance. Although in general the tag cloud functions as a sampling of the more frequent tags within a time span, when transitioning between two time spans tags that are common to both are always included so that the coherence between them can be visualized. Figure 2 describes the animated transition of tags for a single category between two example tag clouds. $t1$ shows the tag cloud for the category during the first time span. At $t2$, tags that are in both sets of tags begin to change size, position, and color transparency. Also at $t2$, tags that are not in the second set of tags begin to shrink. At $t3$ the tags that are not in the second set of tags have disappeared and the tags that were not in the first set of tags begin to appear. $t4$ shows the complete formation of the new tag cloud for the new time span.

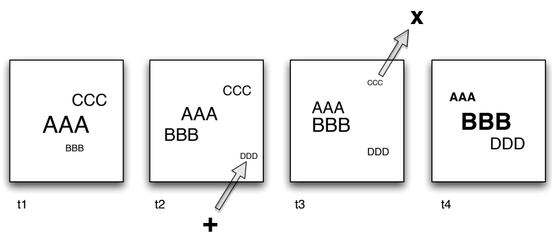

Fig. 2. Diagram of an example transition between two tag clouds. $t1$ shows the initial state of a tag cloud for a particular time span . At $t2$, the A and B tag begin to switch places and the C tag begins to disappear. At $t3$ the C tag has disappeared and a new tag D begins to appear. $t4$ show the new state for the tag cloud at the successive time span. the B and A tag have switched places and are both more opaque, representing that they were both present in the previous time span.

Tag River has two main animation modes. The default mode allows for user interaction. In this mode users can swiftly switch to different time spans and view the detailed tag clouds for each user. A simple mouse rollover triggers the switch, and the general effect is similar to the zooming quality of the Fisheye menu technique [1]. Users can also toggle to a second "ambient" mode in which the time spans shift forward in time at specified increments, looping back to the first time span when necessary. Even during the ambient mode, users can opt to interact with the system. In this "simultaneous" mode, the automatic shifting simply takes place at the new time span the user has interactively moved to. In both of these modes the tag clouds indicate the preservation or dissolution of tags over time. Thus one use of the Tag River visualization is to quickly identify differences in tag clouds over arbitrary times.

3.2 Tag Cloud Layout Algorithm

Bin-packing algorithms attempt to minimize the amount of space taken up by a set of objects. Determining the optimal minimum space is an NP-complete combinatorial problem, and a variety of heuristics have been developed to find good approximate solutions within a minimum amount of time, such as [19]. However, in general these heuristics ignore aesthetic concerns pertinent to information visualization. A popular algorithm to create aesthetically appealing tag clouds is discussed by [20]. And a system that provides flexible control of tag-cloud creation to enhance aesthetic appeal is presented by [12]. We present a customized bin-packing algorithm that handles these issues when displaying the descriptive tags. By providing some "slack" space to each of the tags whereby a tag may be slightly larger or smaller than indicated by its activity we can specify layout considerations, such as text justification, amount of space to fill, arrangement patterns of different size tiles, and quantization of tiles. That is, we seek to find a balance between compelling visual aesthetics and accurate visual communication, particularly in regard to indicating both high-frequency tags within a tag cloud and also tags that persist across multiple time spans.

Our bin-packing algorithm works as follows: 1) the tags are sized and then sorted by their popularity for a particular data region; 2) the first, largest tag is positioned within the centermost space which is large enough to accommodate it; 3) a set of available regions is then identified as possible positions in which to place subsequent tags; 4) a subsequent tag is placed into one of these regions; 5) if it fits then we return to step 3 using the next most popular tag, after appropriately culling the list of possible positions that this placement has invalidated; 6) otherwise we return to step 4 to examine another possible position; 7) if it does not fit anywhere then we find the largest of the possible positions that has a similar aspect ratio as the tag we trying to place and scale it to fit; 8) we then scale all subsequent tags accordingly; 9) we exit the layout algorithm once all tags are placed, or once there are no possible positions in which to place a tag, or once a tag cannot be scaled to fit within any of the possible tags without being shrunk beneath a threshold size.

The most relevant tag is positioned in the center with the largest size, and less relevant tags are positioned further away from the center at smaller sizes. Note that it may not be possible to position the tag in the exact center of the polygon since the top and bottom segments may be slanted. Also, because of the possible scaling of the tags during placement, the sizes of the tags will not necessarily be perfectly representative of their frequency within the dataset. We mitigate this by making all tags with the same or smaller frequency to be similarly scaled so that the general informational aspect of the visualization still functions. Thresholds such as the maximum amount of scaling and the maximum distance of possible positions from the center are controlled by a small set of slack variables. In certain datasets too much or too little slack can adversely affect both the aesthetic display and informational content. As long as we allow non-uniform proportions across the data streams our heuristic is guaranteed to return a solution. If we do enforce a uniform proportionality, then in extreme cases (such as when the quadrilateral has a narrow height or rises at a sharp angle) the layout algorithm may return only a sparse sampling of the tags. In general, the results of this layout method are satisfactory, providing a compromise between aesthetic considerations and spatial accuracy.

4 EXAMPLES

Tag River is a flexible visualization technique that is applicable to a wide variety of data. While the main focus of the Tag River visualization system is to investigate the frequency and temporal coherence of semantic tags, we have explored visualizing mixed sets of images and tags as well as text that does not consist of folksonomic tags. In order to explore its efficacy with social media data, we discuss example data sets that we have used as test cases. These examples utilize the web services APIs from both the Flickr photo-sharing site [5] and the Last.fm online radio site [13]. It should be clear that the Tag River visualization system could be extended to other datasets, as well as to other types of textual data, including the results of other textual analytic tasks. For instance, recent work on topic modeling visualization [8] uses Latent Dirichlet Allocation to automatically extract topics from large unstructured text corpora and then places the topics within an interactive graph. The Tag River visualization emphasizes the temporal aspects of data and thus may be an appropriate alternative visualization for topic modeling (or in fact any textual information) that highlights changes over time.

The examples were written using the Behaviorism [6] framework for data visualization, which provides access to an OpenGL renderer for hardware accelerated graphics and a library of software functionality for gathering and processing data.

4.1 Music Genres and Recording Artists Visualizations

Our first example visualizes user profiles gathered from Last.fm. Last.fm is a music based social networking site that features an online radio station and a personalized music recommendation service. Last.fm builds a detailed profile of each user's musical taste by tracking the details of all the songs the user listens to. Specifically, it provides continually-updated information about users' online presence and also about the songs they have listened to within certain time frames. This information is then publicly accessible through a web services API. We use this information to map musical categories to particular users, and to gather the total time spent online during a particular time frame listening to particular categories. We also extract all the tags associated with genres the user is currently interested in within the specified time frame.

Figure 3 is a detail of the visualization showing the relative frequency of 4 users' listening habits over the course of 30 weeks. Each color represents a different user, and the size of the colored rectangle

Fig. 3. A variation of the Tag river visualization that displays a cloud of music categories for each user during a particular week while simultaneously summarizing the historical data trends.

indicates the the amount of time using the Last.fm site during a particular week (in relation to the other users). Each of the weeks can be expanded to view the most popular genres for each user during that week. The genres take the form of the tag cloud, where size indicates the frequency of a particular genre in relation to the other genres for each user during that week. In Figure 3, week 16 is expanded for detail about the genres, and at a glance, the viewer can discern the general proclivities of each user. For instance, user 3 (third from the bottom) listens to a variety of music within the "metal" and "hard rock" styles, while user 2 (second from the bottom) listens a more eclectic mix of genres ranging from "jazz" to "stoner rock." The Last.fm website will attach multiple genres to a single song, and we make sure that the tag cloud detail of the visualization includes the multiple categories but that the overview trend data does not count any one song more than once.

Our second example also uses the data from Last.fm (and in fact uses the same user profiles and same time period), but uses the tag cloud in the expanded regions to display artist or band names rather than folksonomic tags or categories. In this example, the general trend (indicated by the size of the quadrilaterals for each user) represents the variation of artists that were listened to during each week. More height is assigned to a user's quadrilateral at a particular time step if the have listened to a wider variety of artists during that time. The detail from Figure 1 shows a large amount of fluctuation in this variation for 4 users over a period of weeks. In addition to tracking the general listening patterns of the users over time, the animated tag clouds indicate which artists are listened to week after week, as well as the relative frequency of these artists.

4.2 Multimedia Art Installation Visualization

Our third example experiments with the simultaneous mixing of folksonomic tags with associated images within the tag clouds. Our data is drawn from a dynamic artwork called *Cell Tango* [7] that collects cellphone photography and folksonomic tags describing the photos. Visitors to the exhibition are invited to interactively participate as contributors to the project through the submission of cellphone images. These images become the primary content source of the Cell Tango artwork, and are stored online at the Flickr photo management and hosting website. Cell Tango has been exhibited various fine arts museums and galleries throughout the world for weeks at a time. Thus, Cell Tango provides semi-structured temporal data appropriate for our visualization technique.

Figure 4 shows a detail of the Tag River visualization using the data sets from three different installations of Cell Tango. The first installation (on the bottom) was an exhibition in the Davis Museum at Wellesley College. The second installation (in the middle) was part of a gallery show at the SOMArts Cultural Center in San Francisco. The third installation took place at a media arts festival in Poznan, Poland. Each of these installations occurred at different times, however we have normalized them and arranged them by week so that they can be viewed and compared simultaneously. By showing the activity of each of the installations we engage in a kind of cultural analytics (albeit at

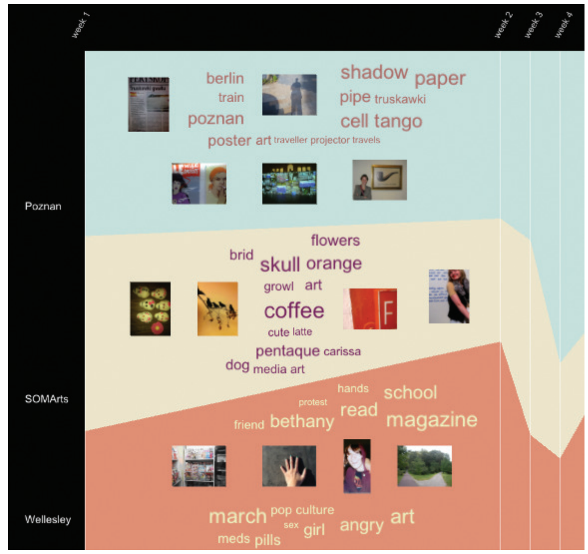

Fig. 4. Detail of the Tag River visualization showing interaction trends of the *Cell Tango* installation at three different galleries.

a smaller scale) such as described by [16], in which aesthetic and/or cultural data points are investigated with various statistical techniques. Since the photos are unique data points, and function more for visual contextualization, the primary analysis is performed on the set of folksonomic tags as they change over time. Scrolling along the weeks spanning the length of the installations provides the user with a sense of the overall topics of interest at the different installation locations. For instance, tags and images from Wellesley reflect a concern with pop-culture, politics, and college. Those from SOMArts on the other hand feature art-related topics, and many of the submitted images are photos of art pieces or of the installation itself. Although the dataset is too small to draw any conclusion from, it seems that the tags used during the first week of the installation in fact influenced both the re-use of specific tags and the general tenor of the user submissions by other participants.

5 Conclusion

Temporal tags are becoming common datasets in many websites utilizing folksonomic categorization. Tag River presents a novel approach that combines temporal data visualization and tag cloud visualization techniques. It uses a bin-packing algorithm that organizes tags within arbitrary polygonal regions based on aesthetic criteria. The Tag River visualization smoothly animates over time, keeping the layout as static as possible, so that both short-term changes and long-term trends are concurrently displayed.

Users currently have only limited interaction possibilities with the example Tag River visualizations. In the future, we plan to extend the interactive capabilities of the application to enable users to determine the ordering and the number of categories, to select multiple time spans, to change the length of the time spans, and to filter unwanted categories and time spans. Additionally, exposing the parameters that control the tag cloud layout and the selection of tags within the tag cloud might be useful. For instance, our tag cloud emphasizes high-frequency tags, but for many folksonomic data sets these high-frequency tags are less interesting and provide less insight than rarer terms. Conducting user studies and creating visualizations to investigate larger datasets and other visual analytic tasks are future steps which will help us to identify and evaluate the strengths and weaknesses of Tag River.

6 Acknowledgments

This work was partially supported by NSF grants IIS-1058132, IIS-0808772, and CNS-0722075.

References

[1] B. Bederson. Fisheye menus. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 217–225. ACM, 2000.

[2] K. Bielenberg and M. Zacher. Groups in social software: Utilizing tagging to integrate individual contexts for social navigation. Master's thesis, Universität Bremen, 2005.

[3] W. Cui, Y. Wu, S. Liu, F. Wei, M. Zhou, and H. Qu. Context-preserving, dynamic word cloud visualization. *Computer Graphics and Applications, IEEE*, 30(6):42–53, 2010.

[4] M. Dubinko, R. Kumar, J. Magnani, J. Novak, P. Raghavan, and A. Tomkins. Visualizing tags over time. In *Proceedings of the 15th international conference on World Wide Web (WWW)*, pages 193–202. ACM, 2006.

[5] Flickr. http://www.flickr.com, accessed September 2011.

[6] A. G. Forbes, T. Höllerer, and G. Legrady. behaviorism: a framework for dynamic data visualization. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization 2010)*, 16(6):101–999, November-December 2010.

[7] A. G. Forbes and G. Legardy. Cell tango: An evolving interactive archive of cellphone photography. In *Proceedings of the Inter-Society for the Electronic Arts Conference (ISEA)*, 2011, to be published.

[8] B. Gretarsson, J. O'Donovan, S. Bostandjiev, T. Höllerer, A. Asuncion, D. Newman, and P. Smyth. Topicnets: Visual analysis of large text corpora with topic modeling. *ACM Transactions on Intelligent Systems and Technology*, 2011, to be published.

[9] Y. Hassan-Montero and V. Herrero-Solana. Improving tag-clouds as visual information retrieval interfaces. In *International Conference on Multidisciplinary Information Sciences and Technologies*, pages 25–28, 2006.

[10] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. Themeriver: Visualizing thematic changes in large document collections. *IEEE transactions on visualization and computer graphics*, pages 9–20, 2002.

[11] O. Kaser and D. Lemire. Tag-cloud drawing: Algorithms for cloud visualization. In *Proceedings of WWW 2007 Workshop on Tagging and Metadata for Social Information Organization*, 2007.

[12] K. Koh, B. Lee, B. Kim, and J. Seo. Maniwordle: Providing flexible control over wordle. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1190–1197, 2010.

[13] Last.fm. http://www.last.fm, accessed September 2011.

[14] B. Lee, N. Riche, A. Karlson, and S. Carpendale. Sparkclouds: Visualizing trends in tag clouds. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1182–1189, 2010.

[15] S. Liu, M. Zhou, S. Pan, W. Qian, W. Cai, and X. Lian. Interactive, topic-based visual text summarization and analysis. In *Proceeding of the 18th ACM conference on Information and knowledge management*, pages 543–552. ACM, 2009.

[16] L. Manovich. Cultural analytics: Analysis and visualization of large cultural data sets. *Screen*, (February):1–23, 2008.

[17] T. Russell. Cloudalicious: Folksonomy over time. In *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries (JCDL)*, pages 364–364. ACM, 2006.

[18] L. Shi, F. Wei, S. Liu, L. Tan, X. Lian, and M. Zhou. Understanding text corpora with multiple facets. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, pages 99–106. IEEE, 2010.

[19] D. Sleator. A 2.5 times optimal algorithm for packing in two dimensions. *Information Processing Letters*, 10(1):37–40, 1980.

[20] F. Viégas, M. Wattenberg, and J. Feinberg. Participatory visualization with wordle. *IEEE Transactions on Visualization and Computer Graphics*, pages 1137–1144, 2009.

[21] F. Wei, S. Liu, Y. Song, S. Pan, M. X. Zhou, W. Qian, L. Shi, L. Tan, and Q. Zhang. Tiara: a visual exploratory text analytic system. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD)*, pages 153–162. ACM, 2010.

[22] Y. Wu, T. Provan, F. Wei, S. Liu, and K. Ma. Semantic-preserving word clouds by seam carving. *Computer Graphics Forum*, 30(3):741–750, 2011.