

Word-Clouds in the Sky: Multi-layer Spatio-temporal Event Visualization from a Geo-parsed Microblog Stream

Masahiko Itoh^{†,‡}Naoki Yoshinaga^{‡*}Masashi Toyoda[‡]

[†] National Institute of Information and Communications Technology, Japan

[‡] Institute of Industrial Science, the University of Tokyo, Japan

E-mail: {imash, ynaga, toyoda}@tkl.iis.u-tokyo.ac.jp

Abstract

Various events, such as public gatherings, traffic accidents, and natural disasters, occur every day in mega-cities. Although understanding such ever-changing events over all these cities is important for urban planning, traffic management, and disaster response, this is quite a huge challenge. This paper proposes a method of visualizing spatio-temporal events with a multi-layered geo-locational word-cloud representation from a geo-parsed microblog stream. Real-time geo-parsing first geo-locates posts in the stream, using geo-tags and mentioned places and facilities as clues. Temporal local events are then identified and represented by a set of words specifically observed in a certain location and time grid, and then displayed above a map as word-clouds. We detect the locality of events to split them into multiple layers to avoid occlusions between local (e.g., music concerts) and global (e.g., earthquakes and marathon races) events. Users can thereby distinguish local from global events, and see their interactions over the layered maps. We demonstrate the effectiveness of our method by applying it to real events extracted from our archive accumulated from five years of Twitter posts.

Keywords— Social media; text analytics; spatio-temporal visualization

1 Introduction

Various sorts of events, such as public gatherings, traffic accidents, and natural disasters, occur every day in mega-cities like Tokyo and New York. These events vary in area, duration, and impact. One disruptive event might cause confusion over wide areas. A key challenge here is understanding these ever-changing events as they happen all over

cities for purposes of urban planning, traffic management, and disaster response.

The rapid growth of social media and smart mobile devices has enabled us to observe the activities and attention of huge numbers of users. Dense and real-time microblog posts are particularly rich spatio-temporal sources of event information. Current public events are almost guaranteed to have microblog users posting what is going on.

Several techniques have been proposed to enable the spatio-temporal events mentioned in such posts visualized on geographical maps overlaid with word-clouds [3, 22] to be better understood. The previous studies, however, have relied on posts explicitly geo-located by users, which have been less than 1% of total posts, and hence, they have missed large numbers of relevant posts. Moreover, they have only visualized single anomalous events and made it difficult to understand temporal changes in areas where multiple events have occurred.

This paper proposes a method of visualizing spatio-temporal events extracted from a geo-located microblog stream as multi-layered geo-locational word-clouds. We first associate posts in the stream with geo-locations by not only using geo-tags but also mentioned places and facilities as clues. The temporal local events are then identified and represented by a set of words specific to a certain location and time, and then displayed on a map. Most events occurred in some specific places (e.g., music concerts), while a few events (e.g., earthquakes and marathon races) affected wider areas. Our method detects the locality of events, and places global events on multiple independent layers to avoid occlusions between local and global events. Users can thereby distinguish local from global events on the map, and see the relationships between them. We demonstrated the effectiveness of our method on real events extracted from our five-year archive of 25-billion Twitter posts.

We outline related work in Section 2 and offer information on our data set in Section 3. We describe our method of visualizing spatio-temporal events in Section 4 and present some case studies in Section 5. Section 6 is the conclusion.

*This work is done while the author was a senior researcher at National Institute of Information and Communications Technology, Japan.

2 Related Work

2.1 Geo-locating Microblog Posts

Existing approaches to spatio-temporal visualization of microblog posts [3, 22] have only exploited posts that are explicitly geo-located by users. The tagged locations usually indicate those of GPS-enabled devices that users utilize to submit posts. Because most users hesitate to expose their locations due to privacy concerns, the location-enabled posts are only a tiny, less than 1% [15], fraction of the entire posts so analysis based on these posts could easily miss posts on local events.

Two distinct approaches have been pursued to cope with the sparsity of explicitly geo-located posts to geo-locate information included in microblog posts, i.e., document geo-location [11] and geo-parsing [13]. Document geo-location estimates the location of authors when they submit posts, while geo-parsing identifies mentions of named entities that are associated with locations (*e.g.*, place and facility names), followed by geo-coding that maps each mention to a specific location. Recent researchers have focused on document geo-location (surveyed in [7]) because explicitly geo-located posts help us to develop (supervised) methods and to evaluate them. The precision of document geo-location is, however, unacceptably low for intra-city level analysis, where the best median error distance is around 30 km, even when using a diverse range of indicators [18]. This is partly because of the task setting of Twitter geo-location. The task estimates the locations of all geo-located posts although they do not always mention experiences related to the locations. In addition, as there are valuable posts that are posted in places irrelevant to user locations, we have adopted the geo-parsing approach in this study.

Watanabe et al. adopted a geo-parsing approach [25] prior to this study, in visualizing spatio-temporal events. They extracted unambiguous facility names from posts submitted through Foursquare¹ and used them to associate posts with their (unique) locations. We extend their approach in this study by propagating the locations of unambiguous facility names to other nouns that co-occur with these facility names to build a large-scale dictionary of geographical entities, and used the dictionary we obtained to associate posts with locations.

2.2 Multi-layer Geo-located Word-cloud Visualization

Some researchers have tackled the detection of abnormal events from Twitter streams and visualized them on maps using word-cloud representations [3, 14, 22, 1]. Most

of them, however, have only extracted global anomalies from explicitly geo-located twitter posts, as was previously stated. Thom et al. [22] extracted clusters of terms based on spatio-temporal locations. However, they were not conscious of relationships between events, and the locality and diffusibility of events. Our approach identifies global and local events from an automatically geo-located microblog stream, and places them on multi-layer word-clouds planes in 3D visualization space.

There have been various sorts of word-cloud layout algorithms [19, 24, 12, 16, 10]. Cui et al. proposed context preserving word-cloud visualizations [5, 26] enabling optimal layouts for word-clouds with multiple timestamps. Although their approaches focused on preserving the positions of words in different layers, they could not consider spatial restrictions. Tag Maps [9] visualized data by placing textual tags on relevant map locations. However, it caused a problem with overlapping texts on the map [21]. Thom et al. [22] laid out texts in a circular way to avoid overlapping them. The circular layout, however, exceedingly broke up the original positions of texts. We thus move texts downward or upward while maintaining their relative positions.

We utilize 2.5D representation for visualizing multi-layers in 3D space. These 2.5D representations are mainly used for visualizing multiple situations in 3D environments such as visualizing different content, visualizing time sequential changes, and visualizing different visual representations and/or models [8, 20, 4, 2]. Although VisLink [4] can link word-cloud layers and map layers, it does not take into consideration geo-locations in preserving multiple word-clouds and their temporal changes.

3 Dataset

We have crawled more than five years' worth of Twitter posts using Twitter API since March 11, 2011. Our crawling started from 30 well-known Japanese users by obtaining their past timelines. We then repeatedly expanded the set of users by following retweets and mentions that appeared in their timelines. We have continuously performed user expansion and tracking of their timelines. Our archive had more than two million users and 25 billion tweets in 2015.

4 Method

Our method of visualizing spatio-temporal events monitors a microblog stream, while detecting temporal local events from the text at a user-defined granularity level of location (grid) and visualizing them as word-clouds on a map, updated at user-specified intervals. Figure 1 outlines the design of the workflow for our method. In what follows, we will describe how temporal local events are found from the stream and visualized them on a map.

¹Foursquare, <https://foursquare.com/>




Figure 1. Overview of our method workflow. We build a gazetteer in advance to find geo-locations of places and facilities mentioned in posts in a microblog stream, and associate the posts with those geo-locations. We next identify temporal local events as sets of words specific to a certain location and time for a user-defined granularity level of location (grid) and time interval. We then detect global events, which affect wider areas, from extracted temporal local events and place them on multiple independent layers. Here, events are represented as spatio word-clouds on each layer.

4.1 Real-time Geo-parsing for Event Detection

Since only 1% of all posts are explicitly geo-located, we associate more posts with locations on the basis of their content. We develop a method that finds posts that can be associated with specific locations. Here, we assume access to a dictionary (or *gazetteer*) that includes pairs of a toponym (or geographical entity) and its geo-locational center. When a post includes a toponym in the dictionary, we associate it with its geo-location. Here, we also assume that the content of the post relates to the toponym’s geo-location [25].

An issue to be addressed is that toponyms could refer to a wide area or several places rather than a specific location. We thus developed a gazetteer of terms strongly correlated to specific locations. The gazetteer was first built from explicitly geo-located posts submitted by location-based services such as Foursquare in our Twitter archive, as in Watanabe et al. [25]. We then fertilized this gazetteer by adding nouns referring only to specific locations. More concretely, we applied a morphological analyzer² to posts to find nouns in the text and associated them with specific locations when they co-occurred with the facility gazetteers. We compute the dispersion of the locations associated with each noun whose frequency was above a threshold. If the variance of the noun’s locations was below a threshold, we added the noun with the mean location to the gazetteer. The resulting gazetteer consists of 38,504 geographical entries.

We use this gazetteer to geo-locate posts in a microblog stream. We adopted a reduced double-array trie [27],³ to

compactly store the gazetteers and enable a quick look-up of the gazetteer entries from text. When posts include gazetteer entries, the posts are associated with the entries’ specific locations. We refer to these posts as *implicitly* geo-located posts to distinguish them from ones explicitly geo-located by geo-tags. We confirm that these implicitly geo-located posts significantly increase the total number of geo-located posts, as we will explain later in Section 5.

As the gazetteer does not include future temporal events in a microblog stream, we identify temporal events from the implicitly and explicitly geo-located posts. We run part-of-speech tagging to find nouns (*terms*), and associate each term w_i with the locations of the posts. Given a user-defined level of location (grid), \mathcal{G} , and time interval, here we compute a TF-IGF (term frequency-inverse grid frequency) score to find temporal events that are specific to each grid $g_j \in \mathcal{G}$:

$$\text{TF-IGF}(w_i, g_j) = \frac{\text{freq}(w_i)}{\sum_{w_i \in g_j} \text{freq}(w_i)} \log \frac{|\{g \in \mathcal{G}\}|}{|\{g \in \mathcal{G} : w_i \in g\}|}.$$

Here, IGF is computed over past intervals to capture bursts of temporal events. Terms with top- n TF-IGF scores are delegated to our visualization engine as temporal local events.

4.2 Multi-layer Spatio Temporal Word-clouds Visualization

The extracted terms are visualized as multi-layer word-clouds on a geographical map. Our visualization engine detects *global events* (terms spread over a wide area), and distinguishes them from *local events* (terms concentrated in specific locations). We use a stacked multi-layer word-clouds design in 3D space to represent these global/local

²MeCab, <http://taku910.github.io/mecab/>
³cedar, <http://www.tkl.iis.u-tokyo.ac.jp/~ynaga/cedar/>




Figure 2. Multi-layer spatio-temporal word-clouds on August 10, 2013 on which two mega-scale events, comic market (Comiket) and fireworks, were held in Tokyo bay area.

events to maintain the spatial relationships between multiple layers (Figure 2). We utilize animation to dynamically display temporal changes in word-clouds on the multiple layers while maintaining spatial contexts. The results look like real clouds of words on the map in the sky.

Small multiples [23] or coordinated multiple views [17] are possible design alternatives for simultaneously comparing multiple events. However, our method generates large indefinite numbers of layers for global events. These designs, which lay out multiple small views in parallel and require a huge screen space for these views, might easily cause small screen problems to display events on huge geographical maps in the small views or to lay out multiple views in the limited screen space. Our stacked multiple layer design enables us to use a large display space for each layer and easily track spatial relationships between events in different layers.

We determine the top- k events to be visualized within an area selected by the user as follows. We first sum the TF-IGF scores of each term in each grid over the past intervals while considering exponential time decay. The resulting score is used to determine the top- k events for the area.

We provide two thresholds that can interactively be defined by the user to identify global events:

Number of clusters: Each occurrence of a term is associated with a specific location (longitude and latitude) of the post. The locations of the term over the past T intervals are clustered by DBSCAN [6],⁴ a density-based spatial clustering. If the number of clusters is greater than the threshold, the term is treated as a global event.

Variance of locations: Some terms have a few clusters but are spread over wide areas. We thus use the variance of the occurrence locations to detect global events.

⁴Users can interactively set two parameters (ϵ and the minimum number of points) for DBSCAN.




Figure 3. Constructing spatio word-clouds.

The threshold is set to six for the number of clusters and 0.006 for the variance of locations in the given examples.

We examine the top-1 to top- k events to check if they are global or not. If a global event is found, a new layer is generated for the event over the base layer, which displays local events. We further check whether the (remaining) top-ranked events (terms) co-occur with the detected global event (term) in the same post. If any, we merge those events into the same layer as the global event and do not examine them later.

Figure 2 shows a multi-layer visualization of the day that has two global events. One event is Comiket (Comic Market), which is the world’s largest amateur comic fair held at Tokyo Big Sight.⁵ Over half a million participants come from all over Japan (and the world), and they wander around Tokyo, talking about Comiket. It is identified as a global event for the entire day, and is displayed on Layer 1 in Figure 2 (I). A fireworks event held on the Harumi pier in the evening is identified as another global event. Both events are displayed on two layers in Figure 2 (II). We can also see local events of “building” and “demolition” on base layer 0 in Figures 2 (I) and (II).

We use word-clouds like representations to track ambulant events such as the Tokyo Marathon or typhoons. As multiple events could appear in almost the same place and overlap each other, it is difficult for us to identify the kinds and sizes of events shown in Figure 3 (ii). We therefore visualize events using term clusters calculated by DBSCAN (Figure 3 (iii)). Each term cluster is represented as a cir-

⁵Comiket, <https://en.wikipedia.org/wiki/Comiket>




Figure 4. Temporal changes in multi-layer spatial word-clouds on February 26, 2012, day on which Tokyo Marathon 2012 was held. Many spectators moved between viewing points to cheer runners while tweeting their situations. Therefore, the sizes and positions of the events in each snapshot are strongly affected by the positions of the runners.

cle, and the term is plotted at its center. The center position of a cluster is calculated as the mean of the occurrence locations weighted by their scores considering time decay. The size of the term and circle are defined by the square root of the total score of their occurrences. We simply resolve conflicts to avoid multiple clusters of events overlapping each other, as shown in Figure 3 (iii), by moving terms with lower scores downward or upward while maintaining their relative positions (Figure 3 (iv)), since the heights of the terms are smaller than their widths in most cases.

Various kinds of word-cloud layout algorithms have been introduced [19, 24, 12]. The layout algorithm used in Wordle [24, 12], which updates positions of words on a spiral of increasing radius to avoid them intersecting with other words, is one of the most fundamental and well-used layout algorithms. However, it breaks up the relative positions of events in a geographical space. We therefore consider that it does not suitable for spatial word-cloud layouts.

We can also optionally plot each occurrence as a small circle to present the dispersion of occurrences because global events are occasionally extracted by variance in the locations of terms. Its size represents its score, while its transparency represents time decay.

Word-clouds on multiple layers often cause occlusion. Users can zoom, rotate, and pan the 3D space to reduce such

occlusion by interactively changing the region being focused on. Moreover, they can interactively control the visibility, height, and transparency of stacked layers to maintain readability and to avoid occlusions. Users can also pan and zoom in/out of maps on this basis. We utilize BingMap⁶ in this implementation to obtain appropriate road maps according to manipulations.

5 Case Studies

5.1 Tokyo Marathon on February 26, 2012

Figure 4 has a visualization of the Tokyo Marathon 2012⁷ held on February 26, 2012 in Tokyo. About 35,000 runners participated, and more than one million people supported or cheered from the roadside.

The runners started in the western part of the course, turned north and then south to the finish (A → B → C → B → D → B → E in Figure 4 (II)).⁸ Most of the (citizen)

⁶BingMap, <http://www.microsoft.com/maps/>

⁷Tokyo Marathon 2012, <http://www.marathon.tokyo/en/info/past/2012/>

⁸Tokyo Marathon course map, <http://www.marathon.tokyo/en/info/course/pdf/map.pdf>

Table 1. Passing time of each point (shown in Figure 4 (II)) for three levels of runners.

Runner level	A	→ B	→ C	→ B	→ D	→ B	→ E
Top runners	9:10	9:40	9:55	10:13	10:34	10:55	11:17
3-hr. runners	9:11	9:53	10:15	10:41	11:10	11:40	12:11
6-hr. runners	9:24	10:49	11:32	12:24	13:22	14:22	15:24




Figure 5. Comparison of visualized events using only explicitly geo-located posts (left) and using both implicitly and explicitly geo-located posts (right).

runners finished the race in three to six hours and the typical times to pass each check point are summarized in Table 1. Many spectators also moved between viewing points to cheer the runners while tweeting their situations.

Figures 4 (I) to (VI) illustrate the huge global event of the marathon, as well as the related events and their temporal evolution. The sizes and positions of the events in each snapshot are strongly affected by the positions of the runners. We can comprehend the detailed situations of the events from the size and positions of related events.

There are 11,927 implicitly geo-located posts via geoparsing, which is significantly larger than that of the 2338 explicitly geo-located posts. Figure 5 compares visualized events extracted from only explicitly geo-located posts (Figure 5 (a)) and both implicitly and explicitly geo-located posts (Figure 5 (b)). We can see that a very limited number of events or no related events are extracted from explicitly geo-located posts.

5.2 Fire at Yurakucho on January 3, 2014

Figure 6 visualizes the evolution of events after a fire broke out near JR Yurakucho station on January 3, 2014, which was during the New Years holidays in Japan when many people visited shrines and temples to make wishes for the coming year. The fire started at around 6:30 am. It caused suspensions of service on the Shinkansen (high-speed inter-city railway) and the JR Yamanote line (loop

line that circles central Tokyo), and hence affected many people’s movements.

The “fire” event in Figure 6 (I) is a local event because people did not start to move so early in the morning. The “fire” in Figure 6 (II) has later become a global event because transportation problems occurred in various places. The “the first shrine/temple visit of the New Year” event on layer 2 in Figure 6 (III) appears at the locations of well-known shrines and temples.

There are 22,482 implicitly geo-located posts via geoparsing, which is again larger than that of the 3070 explicitly geo-located posts.

6 Conclusion

This paper proposed a method of visualizing spatio-temporal events by multi-layered geo-locational word-clouds representation from an automatically geo-located microblog stream. We have demonstrated the effectiveness of our method through an analysis of two temporal events within Tokyo: the Tokyo Marathon and a fire at JR Yurakucho station. We can track the scales and ranges of ambulant and wide-spreading events with our method by using spatio word-clouds, exploring events in detail through observing related events, and observing situations in areas where multiple events occurred using multiple independent layers.

We intend to extend our method to the following direction. We dynamically augment our currently-static location gazetteers to local events that are identified while analyzing microblog streams. We provide functionality that overviews the evolution of temporal events across several time intervals and that displays details of events such as raw tweets or summarized tweets.

Acknowledgements

This work was partially supported by JSPS KAKENHI Grant Number 16K16109 and 16H02905.

References

- [1] Harald Bosch, Dennis Thom, Florian Heimerl, Edwin Puttmann, Steffen Koch, Robert Krüger, Michael Wörner, and Thomas Ertl. ScatterBlogs2: Real-Time




Figure 6. Evolution of events on multiple layers on January 3, 2014, i.e., date on which fire broke out near Yurakucho station. The fire caused suspensions of service on many railway lines and transportation problems in various places. Hence, the “fire” has become a global event. Moreover, as the date was during the New Year’s holidays in Japan, “the first shrine/temple visit of the New Year” event on layer 2 therefore appears at the locations of well-known shrines and temples.

- Monitoring of Microblog Messages through User-Guided Filtering. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2022–2031, 2013.
- [2] Stephen Brooks and Jacqueline L. Whalley. Multi-layer Hybrid Visualizations to Support 3D GIS. *Computers, Environment and Urban Systems*, 32(4):278–292, 2008.
- [3] Junghoon Chae, Dennis Thom, Harald Bosch, Yun Jang, Ross Maciejewski, David S. Ebert, and Thomas Ertl. Spatiotemporal Social Media Analytics for Abnormal Event Detection and Examination using Seasonal-trend Decomposition. In *Proc. VAST*, pages 143–152, 2012.
- [4] Christopher Collins and Sheelagh Cpendale. Vis-Link: Revealing Relationships amongst Visualizations. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1192–1199, 2007.
- [5] Weiwei Cui, Yingcai Wu, Shixia Liu, Furu Wei, Michelle X. Zhou, and Huamin Qu. Context Preserving Dynamic Word Cloud Visualization. In *Proc. PacificVis*, pages 121–128, 2010.
- [6] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proc. KDD*, pages 226–231, 1996.
- [7] Bo Han, Paul Cook, and Timothy Baldwin. Text-based Twitter User Geolocation Prediction. *J. Artif. Int. Res.*, 49(1):451–500, 2014.
- [8] Masahiko Itoh, Naoki Yoshinaga, Masashi Toyoda, and Masaru Kitsuregawa. Analysis and Visualization of Temporal Changes in Bloggers’ Activities and Interests. In *Proc. PacificVis*, pages 57–64, 2012.
- [9] Alexander Jaffe, Mor Naaman, Tamir Tassa, and Marc Davis. Generating Summaries and Visualization for Large Collections of Geo-referenced Photographs. In *Proc. MIR*, pages 89–98, 2006.
- [10] KyungTae Kim, Sungahn Ko, Niklas Elmquist, and David S. Ebert. WordBridge: Using Composite Tag Clouds in Node-Link Diagrams for Visualizing Content and Relations in Text Corpora. In *Proc. HICSS*, pages 1–8, 2011.
- [11] Sheila Kinsella, Vanessa Murdock, and Neil O’Hare. “I’m Eating a Sandwich in Glasgow”: Modeling Locations with Tweets. In *Proc. SMUC*, pages 61–68, 2011.
- [12] Kyle Koh, Bongshin Lee, Bo Hyoung Kim, and Jinwook Seo. ManiWordle: Providing Flexible Control over Wordle. *IEEE Trans. Vis. Comput. Graph.*, 16(6):1190–1197, 2010.
- [13] Jochen L. Leidner and Michael D. Lieberman. Detecting Geographical References in the Form of Place Names and Associated Spatial Natural Language. *ACM SIGSPATIAL Special*, 3(2):5–11, 2011.
- [14] Alan M. MacEachren, Anuj R. Jaiswal, Anthony C. Robinson, Scott Pezanowski, Alexander Savelyev, Prasenjit Mitra, Xiao Zhang, and Justine Blanford.

- SensePlace2: GeoTwitter Analytics Support for Situational Awareness. In *Proc. VAST*, pages 181–190, 2011.
- [15] Stuart Middleton and Lee Middleton. Real-time Crisis Mapping of Natural Disasters using Social Media. *IEEE Intelligent Systems*, 29(2):9–17, 2014.
- [16] Fernando Vieira Paulovich, Franklina Maria Bragion Toledo, Guilherme P. Telles, Rosane Minghim, and Luis Gustavo Nonato. Semantic Wordification of Document Collections. *Comput. Graph. Forum*, 31(3):1145–1153, 2012.
- [17] Jonathan C. Roberts. State of the Art: Coordinated & Multiple Views in Exploratory Visualization. In *Proc. CMV*, pages 61–71, 2007.
- [18] Axel Schulz, Aristotelis Hadjakos, Heiko Paulheim, Johannes Nachtwey, , and Max Mühlhäuser. A Multi-Indicator Approach for Geolocalization of Tweets. In *Proc. ICWSM*, pages 573–582, 2013.
- [19] Christin Seifert, Barbara Kump, Wolfgang Kienreich, Gisela Granitzer, and Michael Granitzer. On the Beauty and Usability of Tag Clouds. In *Proc. IV*, pages 17–25, 2008.
- [20] Zeqian Shen, Michael Ogawa, Soon Tee Teoh, and Kwan-Liu Ma. BiblioViz: a System for Visualizing Bibliography Information. In *Proc. APVis*, pages 93–102, 2006.
- [21] Aidan Slingsby, Jason Dykes, Jo Wood, and Keith Clarke. Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets. In *Proc. IV*, pages 497–504, 2007.
- [22] Dennis Thom, Harald Bosch, Steffen Koch, Michael Wörner, and Thomas Ertl. Spatiotemporal Anomaly Detection through Visual Analysis of Geolocated Twitter Messages. In *Proc. PacificVis*, pages 41–48, 2012.
- [23] Edward Rolf Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1992.
- [24] Fernanda B. Viégas, Martin Wattenberg, and Jonathan Feinberg. Participatory Visualization with Wordle. *IEEE Trans. Vis. Comput. Graph.*, 15(6):1137–1144, 2009.
- [25] Kazufumi Watanabe, Masanao Ochi, Makoto Okabe, and Rikio Onai. Jasmine: A Real-time Local-event Detection System based on Geolocation Information Propagated to Microblogs. In *Proc. CIKM*, pages 1091–1102, 2011.
- [26] Yingcai Wu, Thomas Provan, Furu Wei, Shixia Liu, and Kwan-Liu Ma. Semantic-Preserving Word Clouds by Seam Carving. *Comput. Graph. Forum*, 30(3):741–750, 2011.
- [27] Naoki Yoshinaga and Masaru Kitsuregawa. A Self-adaptive Classifier for Efficient Text-stream Processing. In *Proc. COLING*, pages 1091–1102, 2014.