Spatio-temporal Event Visualization from a Geo-parsed Microblog Stream

Masahiko Itoh
Naoki Yoshinaga
National Institute of Information and Communications Technology
Tokyo 184-8795, Japan
and
Institute of Industrial Science, the University of Tokyo
Tokyo, 153-8505, JAPAN
{imash,ynaga}@tkl.iis.u-tokyo.ac.jp

Masashi Toyoda
Institute of Industrial Science, the University of Tokyo
Tokyo, 153-8505, JAPAN
toyoda@tkl.iis.u-tokyo.ac.jp

IUI'16 Companion, March 07-10, 2016, Sonoma, CA, USA
ACM 978-1-4503-4140-0/16/03.
http://dx.doi.org/10.1145/2876456.2879486

Abstract
We devised a method of visualizing spatio-temporal events extracted from a geo-parsed microblog stream by using a multi-layered geo-locational word-cloud representation. In our method, real-time geo-parsing geo-locates posts in the stream, in order to recognize words appearing on a user-specified location and time grid as temporal local events. The recognized temporal local events (*e.g.*, sports games) are then displayed on a map as multi-layered word-clouds and are then used for finding global events (*e.g.*, earthquakes), in order to avoid occlusions among the local and global events. We showed the effectiveness of our method by testing it on real events extracted from our archive of five years worth of Twitter posts.

Author Keywords
Social media; text analysis; spatio-temporal visualization

ACM Classification Keywords
H.5.2 [Information Interfaces and Representation]: User Interfaces, Graphical user interfaces (GUI); I.2.7 [Artificial Intelligence]: Natural Language Processing, Text analysis

Introduction
In mega-cities like Tokyo, various events, such as public gatherings, traffic accidents, and natural disasters, occur everyday. These events vary in area, duration, and impact.

One disruptive event might cause confusion in a wide area. A key challenge here is understanding these ever-changing events as they happen all over a city 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 of huge numbers of users. In particular, microblog posts are rich spatio-temporal sources of event information. These days, public events are almost guaranteed to have microblog users posting what's going on in it.

To better understand the spatio-temporal events reported in such posts, one can visualize them on geographical maps overlaid with word-clouds [5, 8]. The previous studies however rely on posts explicitly geo-located by users, which is less than 1% of all posts [4], and hence, they miss large amounts of relevant posts. Moreover, they visualize only a single anomalous event, and make is difficult to understand temporal changes in areas where multiple events occur.

We devised 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 using not only geo-tags but also the mentioned places and facilities as clues. The temporal geo-locational events are then recognized as a set of words specific to a certain location and time and displayed on a map. Most of the events occur at some specific places (*e.g.*, concerts), while a few events (*e.g.*, earthquakes) affect wider areas. To avoid occlusions among global events, we detect the locality of events over which to put global events on multiple independent layers. The user can thereby distinguish the local and global events on the map, and also can see relationships between them. We showed the effectiveness of our method on real events extracted from our 5-year archive of 25-billion Twitter posts.

Related Work

Some studies have tried to detect abnormal events from Twitter streams and visualized them on a map [5, 8]. However, they extract only global anomalies from posts explicitly geo-located by users. Instead, we detect local and global events from automatically geo-located Twitter streams and put them on a multi-layer word-clouds planes in a 3D space.

Cui et al. proposed context preserving word-cloud visualizations [2] enabling optimal layouts for word-clouds in multiple timestamps. However, they did not consider any spatial restrictions. Several studies utilized 2.5D representations for visualizing multi-layer in 3D spaces [7, 1]. However, they did not consider multiple geo-location preserving word-clouds and their temporal changes.

Method

Our method monitors a microblog stream, while detecting temporal local events from the text at a user-defined granularity level of location (grid) and visualizes them as word-clouds on a map, updated at user-specified intervals. In what follows, we describe how to find temporal local events from the stream and visualize them on a map.

Since only 1% of all posts are explicitly geo-located, we associate more posts with locations on the basis of their content. Although there are studies on finding a user location from a post, the best median error distance is around 30km [6], which is not accurate enough for analyzing intra-city events. We thus developed a method that finds posts that can be associated with a specific location. Here, we assume access to a dictionary (or *gazetteer*) that includes pairs of a toponym and its geo-locational center. When a post includes a toponym in the dictionary, we associate it with its geo-location (geo-parsing). We assume that the content of the post relates to the toponym's geo-location.

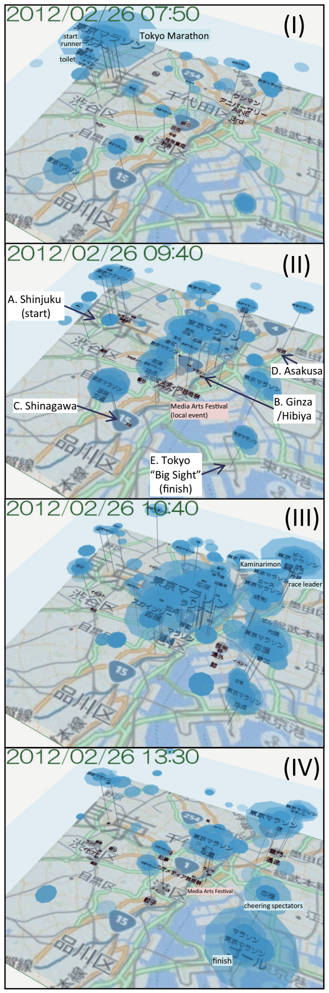

Figure 1: Temporal changes in multi-layer spatial word-clouds on 26 Feb. 2012, the day on which the Tokyo Marathon was held.

An issue here 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 built from explicitly geo-located posts submitted by location-based services such as Foursquare in our Twitter archive. We fertilized the gazetteer by adding named entities referring only to specific locations. Concretely, we used part-of-speech tagging to find proper nouns from the text and associated them with specific locations when they co-occurred with the facility gazetteers. If the variance of the term's locations was below a threshold and the frequency was above a threshold, we added the term with the mean location to the gazetteer. The resulting gazetteer comprised 38,504 entries with specific locations.

We use the dictionary to geo-locate posts in microblog streams. When a post includes gazetteer entries, the post are associated with the entries' specific locations. We refer to these posts as *implicitly* geo-located posts to distinguish them from explicitly geo-located ones using geo-tags.

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

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

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.

The extracted terms are visualized as multi-layer word-clouds on a map. We detect *global events* spread over a wide area to distinguish them from *local events* concentrated in specific grids. To represent these global/local events, we use a stacked multi-layer word-clouds design in 3D space to keep the spatial relationships among multiple layers. We interactively control the visibility and height of the layers to maintain readability and avoid occlusions. We utilize animation to dynamically display temporal changes in word-clouds on the layers while maintaining spatial context. The result looks like real clouds of words on the map (Figure 1).

We determine the top-k events to be visualized within an area selected by the user as follows. It first sums 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 choose the top-k events for the area.

To identify global events, we provide two thresholds that can be interactively defined by the user:
Number of clusters: Each occurrence of a term is associated with a specific location of the post. The locations of the term over the past T intervals are clustered using the DBSCAN algorithm [3]. If the number of clusters is greater than the threshold, the term is treated as a global event.
Variance of locations: There are terms that have a few clusters but are spread over wide areas. We use the variance of the occurrence locations to detect global events.

We examine the top-1 to top-k events 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 if the (remaining) top-ranked events (terms) co-occur with the found global event (term) in the same post. We merge those events, if any, into the same layer as the global event and skip them later.

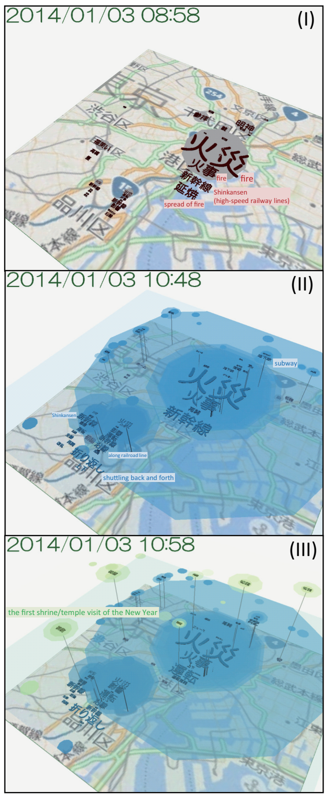

Figure 2: Evolution of events on multiple layers on 3 Jan. 2014, the date on which a fire broke out near Yurakucho station.

To track ambulant events such as marathons or typhoons, we use word-cloud like representations. Each term cluster is represented as a circle, and the term is plotted on its center. The center of a cluster is calculated as the mean of the occurrence locations weighted by their scores considering time decay. The sizes of the term and circle are defined by the square root of the total score of their occurrences. We also plot each occurrence as a small circle to show the dispersion of the occurrences. Its size represents its score, while its transparency represents time decay.

Case Studies

Figure 1 shows a visualization of the Tokyo Marathon held on 26 Feb. 2012. About 35,000 runners participated, and more than one million spectators cheered from the roadside. The runners started in the western part of the course, turned north and then south to the finish (A $\rightarrow$ B $\rightarrow$ C $\rightarrow$ B $\rightarrow$ D $\rightarrow$ B $\rightarrow$ E in Figure 1 (II)). Many spectators moved between viewing points to cheer runners while tweeting their situations.

Figures 1 (I) to (IV) 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 runners. The number of implicitly geo-located posts, 11927, is significantly larger than that of explicitly geo-located posts, 2338.

Figure 2 visualizes the evolution of events after a fire broke out near JR Yurakucho station on 3 Jan. 2014, which was during the new year holidays in Japan when many people visit shrines and temples. The fire started around 6:30 am. It caused suspensions of service on many railway lines and hence affected many peoples' movements. In Figure 2 (I), the "fire" event is still one of local events. Later, in Figure 2 (II), the "fire" has become a global event because transportation problems occurred in various places. In Figure 2 (III), "the first shrine/temple visit of the New Year" event on layer 2 appears at the locations of famous shrines and temples. The number of implicitly geo-located posts, 22482, is larger than that of explicitly geo-located posts, 3070.

REFERENCES

1. C. Collins and S. Carpendale. 2007. VisLink: revealing relationships amongst visualizations. *IEEE Trans. Vis. Comput. Graph.* 13, 6 (2007), 1192–1199.
2. W. Cui, Y. Wu, S. Liu, F. Wei, M. X. Zhou, and H. Qu. 2010. Context Preserving Dynamic Word Cloud Visualization. In *PacificVis*. 121–128.
3. M. Ester, H. Kriegel, J. Sander, and X. Xu. 1996. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *KDD*. 226–231.
4. S. Kinsella, V. Murdock, and N. O'Hare. 2011. "I'M Eating a Sandwich in Glasgow": Modeling Locations with Tweets. In *SMUC*. 61–68.
5. A. M. MacEachren, A. R. Jaiswal, A. C. Robinson, S. Pezanowski, A. Savelyev, P. Mitra, X. Zhang, and J. Blanford. 2011. SensePlace2: GeoTwitter Analytics Support for Situational Awareness. In *VAST*. 181–190.
6. A. Schulz, A. Hadjakos, H. Paulheim, J. Nachtwey, and M. Mühlhäuser. 2013. A Multi-Indicator Approach for Geolocalization of Tweets. In *ICWSM*. 573–582.
7. Z. Shen, M. Ogawa, S. T. Teoh, and K. Ma. 2006. BiblioViz: a System for Visualizing Bibliography Information. In *APVis*. 93–102.
8. D. Thom, H. Bosch, S. Koch, M. Wörner, and T. Ertl. 2012. Spatiotemporal Anomaly Detection through Visual Analysis of Geolocated Twitter Messages. In *PacificVis*. 41–48.