2013 17th International Conference on Information Visualisation

Visualizing Time-Varying Topics via Images and Texts for Inter-Media Analysis

Masahiko ITOH, Masashi TOYODA

Institute of Industrial Science
the University of Tokyo
{*imash,toyoda*}*@tkl.iis.u-tokyo.ac.jp*

Masaru KITSUREGAWA

Institute of Industrial Science, the University of Tokyo
National Institute of Informatics, Japan
kitsure@tkl.iis.u-tokyo.ac.jp

Abstract

This paper proposes a system for analyzing societal behaviors by visualizing time-varying topics in multiple media. Various types of content such as text, images, and videos have spread throughout multiple media, such as TV and the Web, that have complementary information and influence one another. It is important to compare how these media react to real world events to understand recent societal behaviors and how each medium reacts to other media. Our system visualizes flows of content in multiple media in 3D space enabling us to simultaneously explore them. We present two example applications using our system. The first involves the visualization of inter-media events comparing the exposure of topics in TV news and the activities of bloggers. The second example application is a system for visualizing visual trends on social media that chronologically displays extracted clusters of images on blogs. The proposed systems enable users to visually monitor changes in thought, activities, and interests of people, and differences between media through interactively exploring flows of texts and images extracted from the media.

1 Introduction

The mass media has been one of the most popular information sources, and there were few methods to directly measure their influences before social media appeared. Emergence of social media gave us ways to research interests of people from their contents. Our use of media has changed dynamically in the last decade, and has affected our societal behaviors. We now usually watch both mass media, and social media such as blogs, Twitter, and Facebook in our daily lives. Both types of media often influence each other. Many users write articles on social media about topics in TV programs, and TV shows occasionally take up topics that many people discuss on social media. We can therefore investigate influence of mass and social media by complementary analyzing trends in both media. It is important for marketing, politics, and sociology to analyze what kinds of topics have spread through mass or social media, and to analyze differences between multiple media.

Various types of content in media, such as text and images, complementarily help us to understand their ideas, messages, or stories. We can understand activities and their contexts in detail from textual representations. However, it is sometimes difficult to represent an atmosphere and impressions of events by only using text. Even a single photograph might remind us of complex events. Images play a role as effective proxies for content that visually tells us stories of our interests and experiences [1]. There are representative images, such as the design of products and commercial pictures, along with events in the real world.

Our goal in this work is to recognize changes in trends in people's ideas, experiences, and interests through visualizing the flows of texts and image content together, and to interactively analyze their overviews and details. There has been much research that has visualized and analyzed temporal changes in topics on various kinds of media using textural information [22, 14, 4, 3] or images [8, 17, 6, 5]. They have not, however, simultaneously used texts and images to explore trends in media. Crandall et al. [2] and Liu et al. [18] used both textural and image features to extract image clusters or textural tag information; however, the focus of their research was not on extracting temporal changes in trends from texts and images.

This paper introduces a system for visualizing time-varying information complementarily using texts and images extracted from various media resources to analyze society. First, a function to visualize the flow of various images at each timing is required to visually analyze trends from flows of images that describe various topics. We therefore adopt a histogram of images by stacking them on a timeline. This design enables us to find the timing for the beginning of the topic, changes in trends for the topic, burst-

1550-6037/13 $26.00 © 2013 IEEE
DOI 10.1109/IV.2013.77

ing points, and a lifetime of the trends. Second, a function to visualize images on different aspects is necessary to easily compare multiple situations. It is difficult to display an overview of a huge number of images belonging to multiple topics in 2D space at once. We hence arrange multiple histograms of images in 3D space. This design allows us to observe different situations between different topics, sequences of trends, and events with the same timing on different topics. Third, we overlay additional sequences of frequencies extracted from different resources or different aspects as line charts to compare similarities and differences between different media resources. This enables us to find interesting events that are not extracted using only one media resource, such as events that became hot topics in only blogs, or events that became popular on blogs earlier than on TV. Fourth, we integrate an event visualization component using text analysis mechanisms [14] into our system to explore events related to a selected image, topic, and timing in detail. 3D space is necessary to simultaneously show the frequency of images and changes in the structure of events.

We present two example applications in this paper. The first is a visualization system for inter-media events that extracts captured images on TV from a news video archive and extracts bloggers' activities using dependency analysis from a blog archive. It then simultaneously visualizes them in 3D space. The second example application is a visual trends visualizer of social media that extracts and visualizes clusters of images based on visual, textual, and chronological similarities.

2 Related Work

Much work has been done in recent years on visualizing time-varying textural and/or image data to explore trends.

Researchers in the area of time-varying visualization for textual data have developed a number of approaches based on techniques of stacked line charts to visualize changes in trends on topics. ThemeRiver [9] provides methods of visualizing changes in the values of multiple attributes on a timeline. TIARA [22] simultaneously combines ThemeRiver with tag-clouds to visualize changes in keywords consisting of topics. LeadLine [5] provides flow-like metaphors and arranges these flows in parallel to represent topical themes over time.

Our approach adapts multiple line charts or histograms to visualize changes in volumes of information for multiple topics in 3D space. Imoto et al. [10] placed a set of polylines in 3D space to provide two different viewpoints that enabled users to explore both overviews of data from a top view and details of specific parts of data from a frontal view.

Gomi et al. [8] visualized images categorized by time, location, and people in life log data to visualize flows of images. Flake's Pivot [6] provided visualization of cover photographs of magazines from a particular facet. His method used a histogram that displayed images from a particular year from the selected facet. Image Depot [17] visualized flows of images from captured data packets at every IP address to check inappropriate use of the Internet. Compared with the above visualization research on temporal image flows, our system enables us to simultaneously compare histograms of images related to multiple aspects.

There has been some research on extracting clusters of images using visual features [21, 15, 16]. Crandall et al. [2] proposed a method of predicting locations from the visual, textual, and temporal features of photographs that people took of the locations. Liu et al. [18] introduced a tag ranking method using visual and textual features to extract tags related to selected images. These researchers, however, did not mention temporal changes in clusters. Our system provides a method of extracting clusters of images based on visual, textural, and temporal features. Moreover, it chronologically visualizes extracted clusters to explore changes in trends in society.

3 3D Visualization System for Topics via Time-Varying Images and Texts

Our visualization system consists of two main parts such as the Image Flow View and Event View (Figure 1). Users can observe an overview of trends and find interesting events on the Image Flow View, and explore details on the events related to interesting topics and timing on the Event View.

We utilize 3D space to simultaneously represent multiple histograms of images by stacking images and the structures of events on a timeline. The 3D space enables us to simultaneously visualize both flows of images and details on events. Users can zoom, rotate, and pan the 3D space to interactively change the region being focused on and to avoid problems with occlusion.

The proposed system is implemented as a composite component of the IntelligentBox system [20], which is a component-based visual software development system for interactive 3D graphics applications.

3.1 Visualizing Image Flows

Our system visualizes changes in images in topics in 3D visualization space (its basic design has been presented as a poster [12]). It uses the x-axis for the timeline.

We adapt a histogram of images by stacking images on the timeline to represent the flow of various images at each timing. Images in a topic are aggregated per specified time window, such as a month, a week, or a day, and are stacked on the timeline (described as a histogram of images). It uses the y-axis to stack images on the topic with a specified time

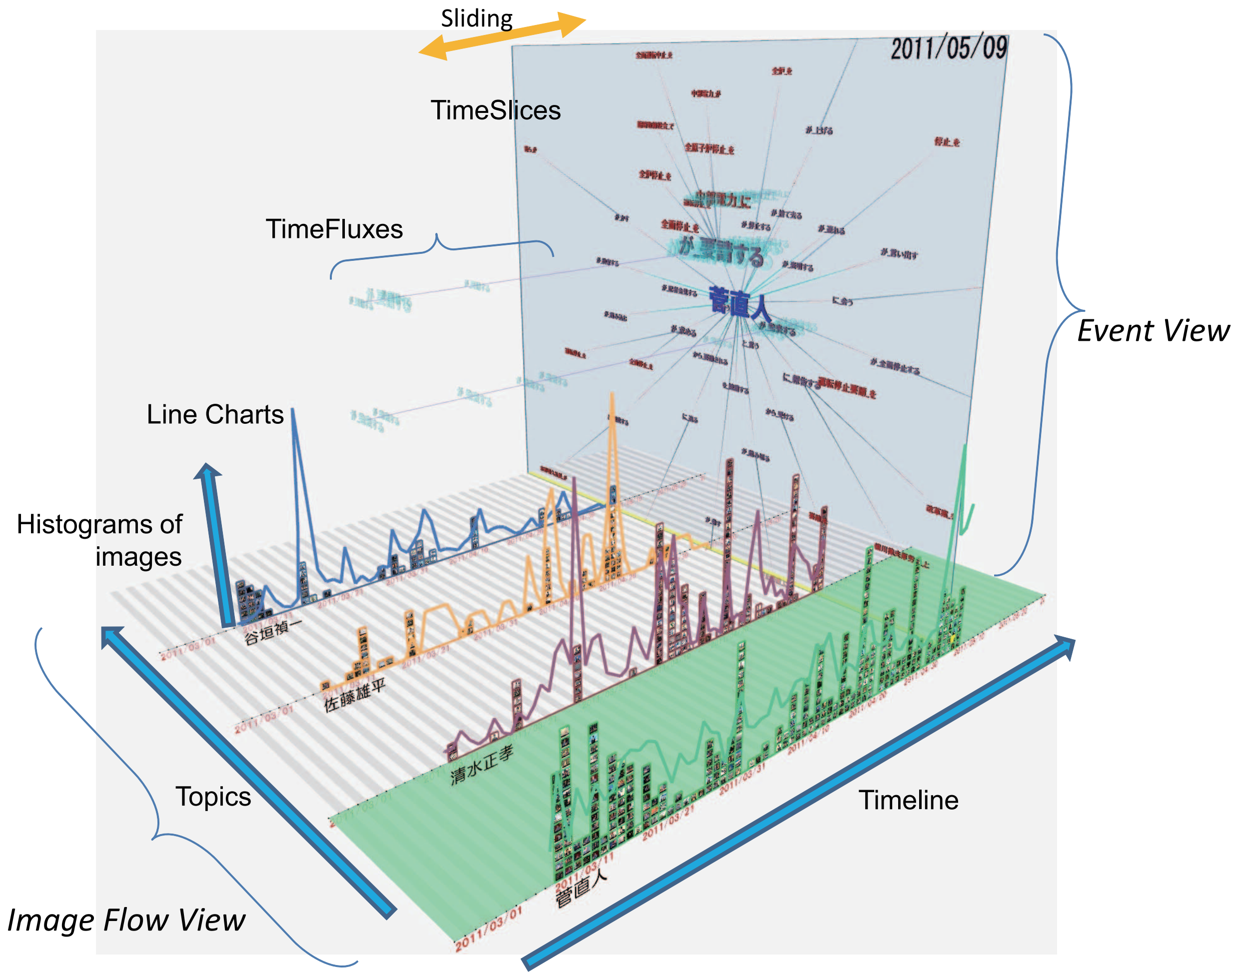

Figure 1. Overview of proposed system to visualize time-varying information complementarily using texts and images. This example is visualization system for inter-media events that displays images from TV on the Image Flow View and bloggers' activities on Event View to analyze society.

window (see Figure 1). It enables us to find the birth timing, bursting points, changes in popular content, and the lifetime of trends for each topic.

We then arrange multiple image flows in the 3D space to compare multiple topics. Topics represented by the histograms of images are arranged along the z-axis in the 3D space (see Figure 1). This allows us to explore differences in bursting timing for every topic, their chronological order, and events with the same timing on different topics.

Users can manually or automatically define the order of topics on the z-axis by using their rankings if they have them. Each topic is colored differently. The frame of each image uses the same color for each topic. The timelines for topics are also colored in the topics' hues. These colored lines represent the active periods of topics (such as those in Figure 5).

Labels that annotate topics can be visualized under the stacks of images. We can present keywords for a topic (Figure 1), or summaries of the topic extracted from textual data (Figure 5, 6) as annotations. Although we can display an arbitrary number of labels, we normally visualize one or two labels from the viewpoint of legibility.

Moreover, users can interactively select images to zoom in on and see them in detail (e.g., Figure 6). The selected images are highlighted and floor panels that have the same time windows as the selected images are also highlighted to support the comparison of multiple topics.

The system supports interaction to explore detailed information on selected images. Users can access original information including a selected image. For example, the blog entry related to the selected image in Figure 6 is opened by the Web browser. The system also has a function to connect images with our text based system of visualizing events (described in Section 3.3) to explore detailed events related to keywords and selected timing.

3.2 Comparing Trends between Different Types of Media

It is possible to provide additional line charts to represent changes in frequencies such as the frequencies of appearance by keywords in other media resources. The line charts are overlaid on the histograms of images to simultaneously compare trends. We use the same time window as the histograms of images to count the frequencies for the line charts. The color schemes for the line charts are the same as the colors of the topics.

3.3 Visualizing Details of Events Analysis

Event View visualizes events, each of which is defined as a set of dependency relations on a particular verb and is regarded as an essential context for a topic, to explore detailed information about a selected topic and timing. When users select a topic and timing, the Event View retrieves events related to the topic and automatically moves to the point of timing on the timeline to display events belonging to the time window. Users can hide it if it is not necessary.

We modify our event visualization component in the proposed system using text analysis mechanisms [14]. It consists of two components such as TimeSlice [13] and TimeFlux [11].

A TimeSlice is a slidable 2D plane on a timeline in the 3D environment, and it visualizes summarized events on the topic keyword during a selected time window as a tree representation (Figure 3). Sliding operation for the TimeSlice along the timeline indicates changes in the structure and frequencies of events. The sizes of nodes represent the frequency of events in the selected time window.

Figure 3 has an example of an event tree for a selected keyword. First, we place a topic keyword at the center of the tree in the TimeSlice and arrange verbs, on which the keyword depends, around the keyword. We then arrange nouns, each of which depends on the verb, as sub-trees of keyword-verb trees. Paths from the keyword to verbs or nouns enable us to construe the context and summary of events. For example, we can understand the event "Prime minister Kan requests shutdown of nuclear power stations to a power company" that occurred on May 7, 2011 from the event tree in Figure 3.

We adopt automatic and dynamic graph layout algorithms based on a force-directed model [7] to visualize the tree structures, which considers a graph (in this case a tree) to be a physical system.

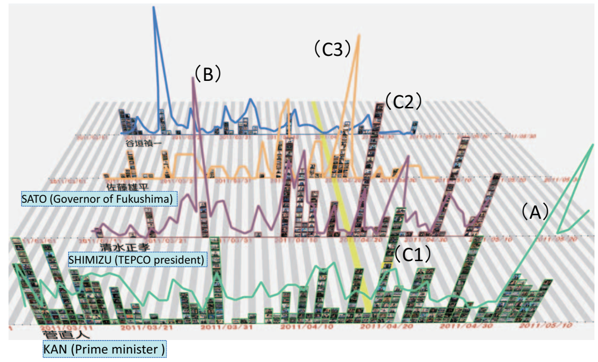

Figure 2. Image Flow View that represent histograms of images from TV corresponding to four key people related to "nuclear power plant", and line charts that represent frequencies of their appearance on blogs to compare different types of media.

A TimeFlux is a line of bubbles, in this case 3D polygon fonts, that visualizes changes in the amount of information such as the number of events within a given period of time (Figure 1). They enable us to intuitively observe the timing when a selected event attracted attention such as a bursting point, and the periodicity of its trends. We can present multiple TimeFluxes by selecting nodes. These allow us to observe differences in trends between other events.

4 Case Study 1: Inter-Media Event Exploration on TV and Blogs

This section explains a first case study using the proposed visualization system that enables users to explore the reactions of society related to particular events such as "an accident at a nuclear power plant" extracted from TV and blog content. First, we introduce a method of extracting images from a news video archive and of extracting summarized event information from a blog archive. We then provide an example to visualize images and events that correspond to key people related to "the accident at a nuclear power plant".

4.1 Extracting Images from TV Archive

News video images related to specified keywords are extracted from a broadcast news video archive created by National Institute of Informatics in Japan.

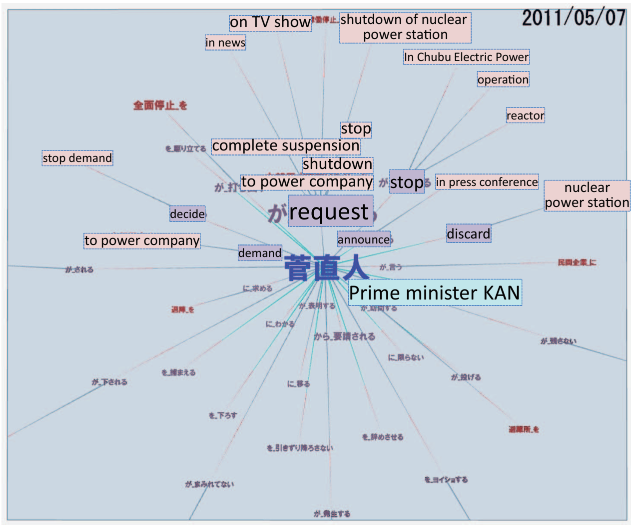

Figure 3. Events on May 7, 2011 related to "Prime Minister Kan" (shown at root of tree structure) extracted from our blog archive. Purple nodes represent verbs, and pink nodes represent nouns.

We first extract a list of keywords from our six-year blog archive including 2 million blogs and 860 million articles. We extract the top 70 names of key people related to a "nuclear power plant" in this example from the blog archive by using named entity extraction and count them. Second, we extract a list of dates, times, and names of broadcasting stations, where names of people were mentioned on the news on TV, from the closed captioning data of the news. We then extract news video frames that correspond to each date, time, and name of the broadcasting station from the news video archive from March 1, 2011 to June 30, 2011.

4.2 Extracting Events from Blog Archive

We generate an event database related to these key people from the blog archive. We first extract sentences in which the names of key people appeared from the blog archive for this purpose. We then extract phrase dependency structures, in which each dependency relation includes the key people, from the extracted sentences by using dependency analysis, and then construct the event database. Events are aggregated per day in this example. Details on the event database are provided in Itoh et al. [14].

Moreover, we extract the frequencies of events related to people every day for the line charts to compare trends in blogs with those on TV (Figures 1 and 2).

Figure 4. Peaks of histograms for images common to three people. People appear in different image flows. (This snapshot filters out other images.)

4.3 Example of Visualization

Figures 1 and 2 indicate the number of mentions on TV and blogs every day corresponding to the four key people after "the Great East Japan Earthquake"; such as Kan (Prime minister), Shimizu (president of Tokyo Electric Power Company (TEPCO)), Sato (governor of Fukusima prefecture), and Tanigaki (president of the Liberal Democratic Party in Japan) from front to back. The topics are manually ordered by users, which enables them to explore the reactions of society to the activities of each individual after the disaster from the aspects of both mass and consumer generated media.

There is a long peak in the line chart on the blog about Prime minister Kan in Figure 2 (A) compared with the histogram of images on TV. Figure 3 visualizes detailed events that represent the activities of Prime minister Kan and people's thoughts and/or interests in his actions on May 7, 2011. We can see that when Prime Minister Kan asked electric power company to stop generating nuclear power, it became a huge topic on the blog. Many people discussed

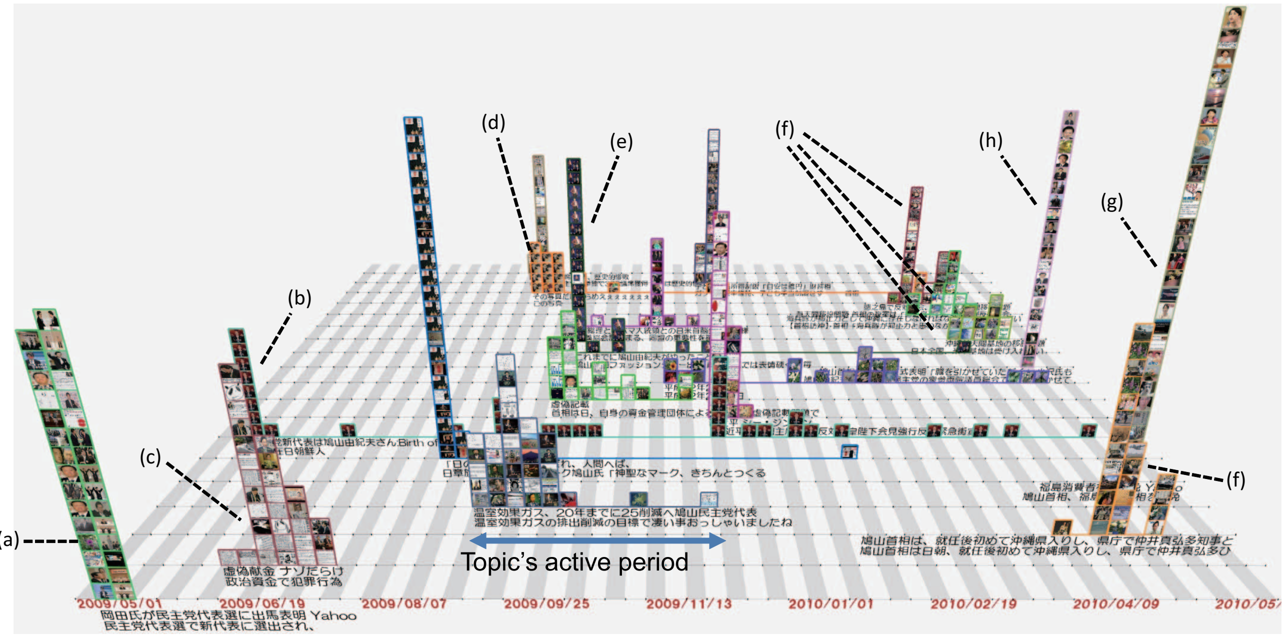

Figure 5. Example visualizing changes in clustered images for given query related to “Prime Minister Hatoyama”.

this topic on the blog because it was very sensitive. We can also see that Prime Minister Kan continued asking power company about various matters after May 7, 2011 through the TimeFluxes in Figure 1.

We can see that the topic related to president Shimizu became a hot topic on the blog earlier than on TV from Figure 2 (B). We can see, by using Event View, that many bloggers criticized him for not appearing after the nuclear accident, then he was hospitalized. TV only reported when he was hospitalized.

Figures 2 (C1), (C2), and (C3) indicate there are peaks for the histograms of images common to the three people (Kan, Shimizu, and Sato). We can also see that they appear in different image flows from Figure 4. This means that they are simultaneously concerned with the same events. The frames with the same colors represent the same or almost the same images in Figure 4. Prime Minister Kan visited Fukushima prefecture to meet its governor in this example, and the president of TEPCO also visited him there. These peaks on the histograms of images indicate that images of these events are excellent materials for content on TV shows.

Our visualization system enables users to recognize important events in this way from the Image Flow View, and then explore their details by using the Event View.

5 Case Study 2: Visual Trends in Social Media

This section explains visual trends in the blogosphere caused by images in articles. Blog articles include various images such as photographs of real-world events, images of bloggers’ favorite goods, and illustrations explaining the details of news. We can find trends through flows of typical images that describe the events by organizing these images, and can observe changes in bloggers’ interest in news and consumer products.

5.1 Extracting Visual Trends from Blog Archive

Our system retrieves relevant articles from our blog archive for a given query, and then extracts images and surrounding texts included in the articles. These images are clustered based on their visual, textual, and chronological similarities, and then visualized as an image flow.

We adopt local descriptors as visual features that are relatively robust to scaling, rotation, and affine transformations. Typical images in blogs are photographs taken by bloggers, and they often include the same object taken from different angles. Other typical images are copies that have been slightly modified from different media. Local descriptors are suitable for capturing similarities in such images. We use Lowe's implementation of SIFT (Scale Invaliant Feature Transform) features [19] in this work. The similarities between a pair of images is calculated by the number of matching feature points divided by the average number of feature points in both images.

We use term frequency vectors as textual features. We extract several lines of text surrounding each image. The similarities between texts are calculated by the cosine of term frequency vectors weighted by tf-idf.

Both similarities are linearly combined and decayed over time, since images of the same topic tended to appear at around the same time. The following equation derives the final similarity:

$$S(A,B) = \gamma \cdot S_{image}(A,B) \cdot e^{-\alpha t} + (1-\gamma) \cdot S_{text}(A,B) \cdot e^{-\beta t},$$

where t is the difference in the timestamps of images A and B, and α, β, and γ are the empirically determined parameters.

Images are clustered into sub-topics with a method of agglomerative hierarchical clustering using the similarity measure. Then, each cluster is labeled by using sentences chosen from the surrounding text that includes terms that frequently appear in the cluster. Clusters are ranked based on the number of images and similarities between images in each cluster. Finally, the results are visualized with our proposed system.

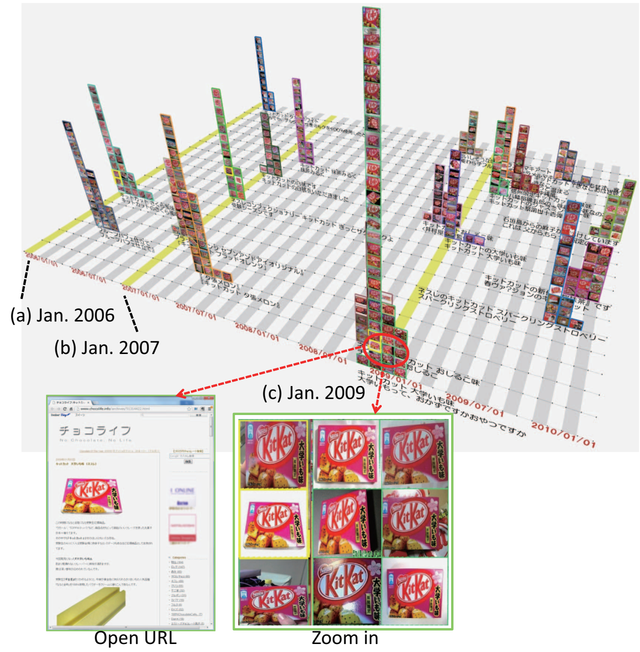

Figure 6. Visualizing changes in trends related to varieties of "Kitkat"

5.2 Examples of Visualization

5.2.1 Visual Trends on Topics about Prime Minister Hatoyama

Figure 5 visualizes clustered images for a given query related to "Prime Minister Hatoyama", where the top 20 clusters are arranged from front to back according to their rankings. Images are aggregated per week. We can read stories about "Prime Minister Hatoyama" by exploring the movements of topics.

(a) Hatoyama was elected as president of the Democratic Party of Japan at an election on May 16, 2009 to replace the previous president, Ozawa. Hatoyama took office as Prime Minister just after that. (b) A collage image representing "Hatoyama is a marionette of Ozawa" became popular. This image alluded to many people who considered Ozawa to be a fixer for the Hatoyama government. This cluster appears repeatedly when various problems occurred in the Hatoyama government, and continues until it was ousted. (c) A cluster related to a false contribution scandal appears. (d) Strange photographs of his youth spread. (e) Prime Minister Hatoyama appeared on a fashion show. (Clusters (c) and (d) became popular only on the Web though social media). (f) The problem with the relocation of U.S. military bases on the island of Okinawa became huge (there are four clusters related to this problem). After event (f), the image cluster for the dismissal of a cabinet minister (g) becomes enormous because she was opposed to the Prime Minister about the problem with bases on Okinawa. Finally, Hatoyama resigned at (h) just after event (g).

The result shows the differences between political topics on Social Media; Events that caused criticism or argument such as (c) or (f) remain for a long term on the timeline, while events that were reported a lot on news and did not cause argument such as (h) and (g) remain for a short term. Satirical images such as (b) and (d) tend to be reused many times.

5.2.2 Visual Trends on Flavors of Kitkat

Figure 6 visualizes changes in trends of new products related to "Kitkat" confectionery in Japan. The clusters of images represent the different kinds of flavors for "Kitkat". They are sorted by cluster rankings. Images are aggregated per month. We can find seasons when various kinds of new flavors for "Kitkat" become popular, especially in January.

"Cherry Kitkat" became popular in January 2006 and 2007 (as seen in Figures 6 (a-b)) because students " prayed for a pass in their entrance exams" since the sentence "a cherry tree is blooming" is sometimes used to represent "success in an examination" in Japan and "Kitkat" means "surely win" in a play on words in Japanese. The cluster related to "candied sweet potato Kitkat" became popular in January 2009 in Figure 6 (c) for almost the same reason. Users can browse blog entries including selected images to check the details of contexts for the reasons many people mentioned this flavor during this month (Figure 6).

The result shows that kitkat not only having good taste but also having meaning or wish tend to become popular on the Web.

Conclusion

We proposed a visualization system to explore trends and events in various types of media such as TV and/or blogs through observing the flows of images and changes in event tree structures. We used images extracted from a video archive on broadcast news and a blog archive in the case studies. Our approach could also be applied to images from other types of media such as Twitter and Flickr. Events were not only extracted from blogs but also from closed captioning and Twitter for comparison.

Our future work includes expanding our current work that allows users to easily recognize the characteristic of events from image flows using statistical values such as the mean and variance of histograms to extract important events, and cross correlation histograms from different types of media to extract lead lags from them.

Acknowledgements

We would like to thank Professor Shin'ichi Satoh (National Institute of Informatics) for providing us with invaluable data from a video archive of broadcast news and Associate Professor Naoki Yoshinaga for constructing event databases and services using dependency analysis. We also wish to thank Mr. Tetsuya Kamijo for the results of image clustering described in Section 5. Further, we would like to thank Associate Professor Norio Katayama (National Institute of Informatics) for extracting and providing video images from the news video archive, and Associate Professor Sayaka Akioka (Waseda University) for extracting and providing the list of timing when the names of key people were mentioned in the closed captioning data described in Section 4. This work was supported by the Multimedia Web Analysis Framework toward the Development of Social Analysis Software Programs by the Ministry of Education, Culture, Sports, Science and Technology, Japan.

References

[1] Amy Campbell, Christopher Wienberg, and Andrew S. Gordon. Collecting Relevance Feedback on Titles and Photographs in Weblog Posts. In *Proc. of IUI '12*, pages 139–148, 2012.

[2] David Crandall, Lars Backstrom, Daniel Huttenlocher, and Jon Kleinberg. Mapping the World's Photos. In *Proc. of WWW '09*, pages 761–770, 2009.

[3] Weiwei Cui, Shixia Liu, Li Tan, Conglei Shi, Yangqiu Song, Zekai Gao, Huamin Qu, and Xin Tong. TextFlow: Towards Better Understanding of Evolving Topics in Text. *IEEE Trans Vis Comput Graph*, 17:2412–2421, 2011.

[4] Weiwei Cui, Yingcai Wu, Shixia Liu, Furu Wei, Michelle Zhou, and Huamin Qu. Context-Preserving, Dynamic Word Cloud Visualization. *IEEE Computer Graphics and Applications*, 30:42–53, 2010.

[5] Wenwen Dou, Xiaoyu Wang, Drew Skau, William Ribarsky, and Michelle Zhou. LeadLine: Interactive Visual Analysis of Text Data through Event Identification and Exploration. In *Proc. VAST '12*, pages 93–102, 2012.

[6] Gary Flake. Is Pivot a Turning Point for Web Exploration? In *http://www.ted.com/talks/gary_flake_is_pivot_a_turning_point_for_web_exploration.html (TED.com video - Feb/Mar 2010)*.

[7] Thomas M. J. Fruchterman and Edward M. Reingold. Graph Drawing by Force-Directed Placement. *Software Practice and Experience*, 21(11):1129–1164, 1991.

[8] Ai Gomi and Takayuki Itoh. A Personal Photograph Browser for Life Log Analysis based on Location, Time, and Person. In *Proc. of SAC '11*, pages 1245–1251, 2011.

[9] Susan Havre, Elizabeth Hetzler, Paul Whitney, and Lucy Nowell. ThemeRiver: Visualizing Thematic Changes in Large Document Collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):9–20, 2002.

[10] Maiko Imoto and Takayuki Itoh. A 3D Visualization Technique for Large Scale Time-Varying Data. In *Proc. IV '10*, pages 17–22, 2010.

[11] Masahiko Itoh. 3D Techniques for Visualizing Users' Activities on Microblogs. In *Proc. the IET International Conference on Frontier Computing 2010*, pages 384–389, 2010.

[12] Masahiko Itoh, Masashi Toyoda, Tetsuya KAMIJO, and Masaru Kitsuregawa. Visualizing Flows of Images in Social Media. In *Proc. of VAST '12 (poster)*, pages 129–130, 2012.

[13] Masahiko Itoh, Masashi Toyoda, and Masaru Kitsuregawa. An Interactive Visualization Framework for Time-Series of Web Graphs in a 3D Environment. In *Proc. IV 2010*, pages 54–60, 2010.

[14] Masahiko Itoh, Naoki Yoshinaga, Masashi Toyoda, and Masaru Kitsuregawa. Analysis and Visualization of Temporal Changes in Bloggers' Activities and Interests. In *Proc. of PVis '12*, pages 57–64, 2012.

[15] Yushi Jing and Shumeet Baluja. Pagerank for Product Image Search. In *Proc. WWW '08*, pages 307–316, 2008.

[16] Lyndon Kennedy and Mor Naaman. Generating Diverse and Representative Image Search Results for Landmarks. In *Proc. of WWW '08*, pages 297–306, 2008.

[17] Ko Kubota, Hideki Koike, and Michiaki Yasumura. Image Depot: Research on Image Gathering and Displaying by Packet Capturing. In *Proc. of WISS '09*, pages 167–168, 2009. (in Japanese).

[18] Dong Liu, Xian-Sheng Hua, Linjun Yang, Meng Wang, and Hong-Jiang Zhang. Tag Ranking. In *Proc. of WWW '09*, pages 351–360, 2009.

[19] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vision*, 60(2):91–110, 2004.

[20] Yoshihiro Okada and Yuzuru Tanaka. IntelligentBox: A Constructive Visual Software Development System for Interactive 3D Graphic Applications. In *Proc. CA 1995*, pages 114–125, 1995.

[21] Reinier H. van Leuken, Lluis Garcia, Ximena Olivares, and Roelof van Zwol. Visual Diversification of Image Search Results. In *Proc. of WWW '09*, pages 341–350, 2009.

[22] Furu Wei, Shixia Liu, Yangqiu Song, Shimie Pan, Michelle X. Zhou, Weihong Qian, Lei Shi, Li Tan, and Qiang Zhang. TIARA: A Visual Exploratory Text Analytic System. In *Proc. KDD '10*, pages 153–162, 2010.