KD-Photomap: Exploring Photographs in Space and Time

Iulian Peca Haolin Zhi Katerina Vrotsou Natalia Andrienko Gennady Andrienko

University of Bonn and Fraunhofer Institute for Intelligent Analysis and Information Systems (IAIS)

ABSTRACT

KD-photomap is a web-based visual analytics system for browsing collections of geotagged Flickr photographs in search of interesting pictures, places, and events. Spatial filtering of the data is performed through zooming, moving or searching along the map. Temporal filtering is possible through defining time windows using interactive histograms and calendar controls. Information about the number and spatiotemporal distribution of photos captured in an explored area is continuously provided using various visual cues.

1 INTRODUCTION

Browsing through publicly shared geotagged photographs has become common practice for exploring destinations around the world. Popular web mapping sites like *Google Maps* provide access to geotagged photographs. Photo sharing sites like *Flickr* and *Panoramio* include options for map-based exploration. Even though great browsing flexibility is offered through these interfaces, they do not provide room for deeper analysis of the data or area.

To this end, a number of applications and research initiatives have been concerned with improving the exploration and analysis of geotagged information [6]. Interactive map-based interfaces are becoming the standard way of displaying and spatially filtering the data. To improve the browsing experience, focus has been put in summarizing the available information as representative pictures or tags (eg. [4]) and in filtering photographs by topic [1]. Filtering with respect to both time and topic is addressed in [3] through combining spatial, temporal and 'topical' interactive visualizations. In most cases, however, the primary focus is browsing the data content, and therefore information about the extent and spatiotemporal distribution of the data is commonly insufficient and only relatively displayed. An exception example is BOOMsys [5] which is a visual analytics system for exploring the spatiotemporal distribution of improvised explosive device (IED) incidents.

KD-photomap (KD: Knowledge Discovery) is a web-based application for map-based exploration of Flickr photographs. It provides both an interface for browsing photos of a region but also a framework for analysing the spatiotemporal characteristics of the region itself based on this data. Common geobrowsing principles are combined with visual analytics tools to allow a user to get an overview of the spatial distribution of the data as well as detect, explore and compare spatial and temporal patterns across the map.

2 KD-PHOTOMAP

KD-photomap is composed of an Oracle database containing Flickr photograph metadata, which is accessed through a client-server structure. The system is developed using Java, JavaScript, Google Maps API v3, Yahoo! API, and the Flickr API for data collection.

2.1 User Interface

The KD-photomap user interface viewed on the web-browser (client) is composed of four main elements (fig. 1): (1) the background *Google map*, (2) a *colour legend* panel, (3) a *control panel*, and (4) the *'carousel'* which is the photo display panel.

IEEE Symposium on Visual Analytics Science and Technology
October 23 - 28, Providence, Rhode Island, USA
978-1-4673-0013-1/11/$26.00 ©2011 IEEE

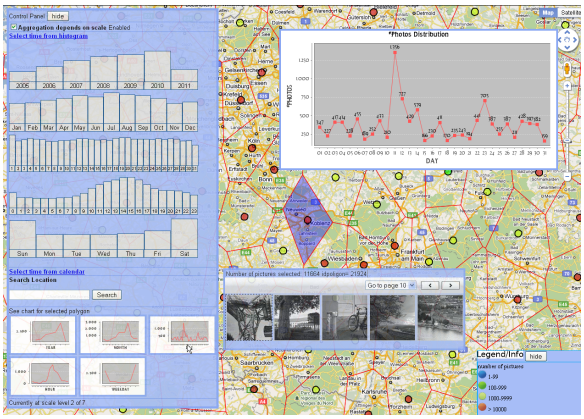

Figure 1: Overview of KD-photomap interface.

KD-photomap is built around the *Google Map* and retains its original Google controllers. In addition, the displayed map region is divided into cells, defined by Voronoi polygons, reflecting the spatial density of the captured photographs. The tessellation is performed using the method described in [2] and is displayed through a red polygon layer. A coloured circle is drawn in the centre of each cell reflecting the total number of photographs there. The *colour legend* (fig. 1, bottom right) displays the correspondence between the number of photos and the coloured circles.

The *control panel* provides the user with several interaction options and includes five subcomponents. *"Aggregation"*, provides an option to control the size of the cell defined areas. If enabled, the granularity of the space division will be adapted to the zoom level, if not, the space division obtained on start-up will remain unchanged in any zoom level. *"Select time from histogram"*, contains histograms for displaying the number of photos captured in the explored map region with respect to different time resolutions (fig. 2(a)). The distribution of photos over years, months, days, hours and week days between 2005 and 2011 is displayed. Moreover, complex filtering criteria can be interactively defined through the histograms by selecting values at several time resolutions. Hovering the mouse over a histogram bar displays the number of pictures corresponding to that bar's time point, while clicking on it adds this time point to the filtering selections. *"Select time from calendar"*, allows selections of time intervals or points from a calendar, while *"Search Location"* allows a user to 'move' to any place fast by typing a destination in a search bar. *"See chart for selected cell"*, is activated when a cell is selected. Time-graphs are then computed showing the number of pictures in the cell with respect to years, months, days, hours and week days. They are displayed as icons in the control panel and are enlarged when clicked (fig. 1).

The *'carousel'* is the panel for displaying the actual Flickr photographs. It is activated when a cell is selected and contains pictures from the selected cell. Photos are loaded into the carousel 30 at a time and 5 are displayed (fig. 2(b)). On mouse-over the pictures are displayed enlarged and a marker on the map appears, indicating where the picture has been captured. When the user clicks on a picture a new tab showing the Flickr link of the picture is opened in the web-browser.

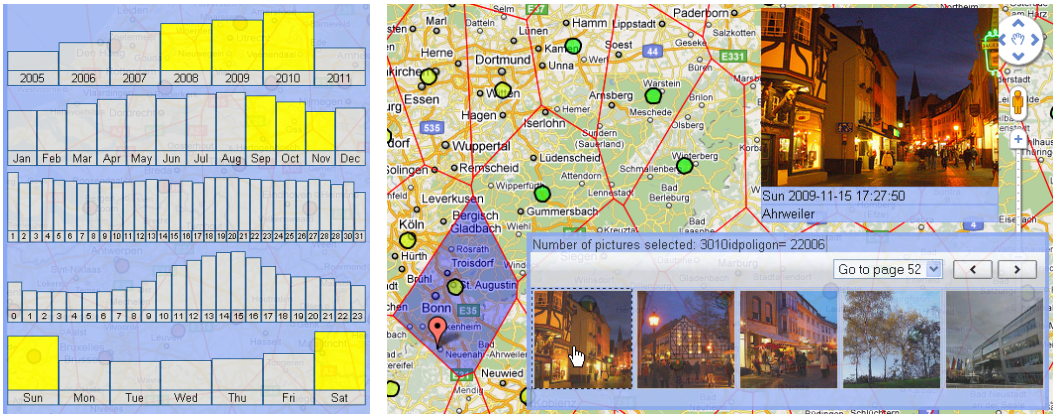

Figure 2: Basic elements of KD-photomap. Left: Interactive histograms provide overview of the temporal distribution of photos captured in the explored map region and are used for creating complex temporal filtering queries. Right: The explored map region is divided into polygon cells based on the spatial density of the photos. Selecting a cell allows a user to browse photos corresponding to it in the display carousel.

2.2 Database

Our dataset is collected from Flickr using the Flickr API. A crawling approach has been used for this, since the API allows the download of photos and contact lists of single users but not of all available data at once. During this crawl a set of users are selected as random seeds and a breadth-first search is performed to recursively explore all their contacts. Only metadata about the shared photographs of each user are saved to the database. These consist of user and photo id, text tag, date of capture, geographical coordinates, and links to the actual photographs which are used to access them. In order to keep the dataset updated, since Flickr does not provide methods for obtaining differences from a given timestamp, the breadth-first search needs to be repeated, which takes approximately 50 days on a standard PC. Until now, June 2011, our collected data include approximately 10 million users, and 110 million geotagged photos in the world and the process is still ongoing.

Moreover, to increase time efficiency in KD-photomap, the Voronoi polygons that divide the explored map region into cells, as discussed in sec. 2.1, have been pre-computed and stored in the database. Geographical coordinates defining these polygons are saved with 7 levels of granularity depending on map zoom level.

Finally, KD-photomap allows users to flexibly create complex spatiotemporal queries. As filtering conditions increase in complexity, the processing time of the queries also increases. In order to optimize this, the number of photos per cell have been precomputed for all possible time interval combinations and stored as long comma-separated strings in the table of cells. Having done this, to calculate the number of photos for each cell, the server just needs to load the cell's time interval string, and sum up all the photo-counts corresponding to the query specifications. Such calculation for each cell takes approximately 50ms on a standard PC, and 5-10s for each query.

2.3 Client-server Communication

KD-photomap is a web-based application built using a client-server structure, therefore, as a user interacts with the client side, requests are created that are sent to the server for processing. All communication between client and server is performed through XML messages. Each XML query message has two main components for all the scenarios, a header and a body. In the header the geographical coordinate boundaries and the centre of the map region displayed on the web-browser are saved. The header remains unchanged regardless of the user options. The body of the XML message includes the temporal filtering specifications, if any. This component changes according to the selections made by the user on the screen. Through this structure of the message we maintain a permanent coordination and control of the information flow received by the user.

After the map has been loaded on the client side and information about its boundaries is established, the displayed map is divided into cells defined by Voronoi polygons. Since, as we discussed previously, these polygons are pre-computed, this is performed by retrieving the set of polygons corresponding to the geographical region and zoom level from the database.

3 Conclusions

KD-photomap provides a flexible interface for geobrowsing photographs and a platform for analysing the spatial and temporal characteristics of photo-capturing activity across the map. The system has several advantages. Overview information about all the data present in the map region under exploration is displayed at all times. Detailed information of selected areas is displayed on demand making it possible to compare them with other areas and with the larger displayed region. Temporal patterns can be detected through the time-graphs displaying the temporal distribution of photos per area. The patterns can give rise to hypotheses that can be investigated further, for example periodical peaks in an area can indicate reoccurring events happening there. In the future we plan to investigate approaches for space-time online event detection [2], and also experiment with additional representations such as tag clouds of cells to get a contextual overview of them. KD-photomap will soon be available for public use.

References

[1] S. Ahern, M. Naaman, R. Nair, and J. Yang. World Explorer: Visualizing Aggregate Data from Unstructured Text in Geo-Referenced Collections. In *Joint Conf. on Digital Libraries*, Vancouver, Canada, 2007.

[2] G. Andrienko, N. Andrienko, M. Mladenov, M. Mock, and C. Pölitz. Discovering Bits of Place Histories from People's Activity Traces. In *IEEE Symposium on Visual Analytics Science and Technology*, pages 59–66, Salt Lake City, Utah, USA, 2010.

[3] M. Dörk, S. Carpendale, C. Collins, and C. Williamson. VisGets: coordinated visualizations for web-based information exploration and discovery. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1205–12, 2008.

[4] A. Jaffe, M. Naaman, T. Tassa, and M. Davis. Generating summaries and visualization for large collections of geo-referenced photographs. In *8th ACM International Workshop on Multimedia Information Retrieval*, page 89, New York, New York, USA, 2006. ACM Press.

[5] M. Murdock, B. Finch, N. Maziekas, and R. Roth. The Basic Ordnance Observational Management System (BOOMsys), 2010.

[6] Y.-T. Zheng, Z.-J. Zha, and T.-S. Chua. Research and applications on georeferenced multimedia: a survey. *Multimedia Tools and Applications*, 51(1):77–98, Nov. 2010.