

COMPARATIVE VISUALIZATION OF GEOSPATIAL-TEMPORAL DATA

Stefan Jänicke^{1,2,3}, Christian Heine¹, Ralf Stockmann² and Gerik Scheuermann¹

¹*Image and Signal Processing Group, Institute for Computer Science, University of Leipzig, Leipzig, Germany*

²*Göttingen State and University Library, University of Göttingen, Göttingen, Germany*

³*Fraunhofer Institute for Intelligent Analysis and Information Systems IAIS, Sankt Augustin, Germany*

Keywords: Visual Data Exploration, Geovisualization, Comparative Visualization.

Abstract: The amount of online data annotated with geospatial and temporal metadata has grown rapidly in the recent years. Social networks like Flickr and Twitter are popular providers of masses of such data, but are hard to browse. Many systems exist that can show the data's geospatial distribution in a map widget and their temporal distribution in a time widget, allowing these widgets to become dynamic-query-like filters for data. We present a web application that combines several existing approaches and supports comparison of multiple, potentially large result sets of textual queries in a geospatial and temporal context with extended interaction capabilities. We validate our approach with several case studies and our tool's learnability in a field experiment.

1 INTRODUCTION

Although the amount and types of data available through public web resources is seemingly endless, the main method of finding information is still largely performed by text queries. Because of the typically enormous amount of results for text queries, popular search engines rank their results using a blend of relevance and popularity. The results often include items the user is not interested in, requiring him to refine the query by adding additional or using different keywords, but repeated refinement can lead to frustration.

While the amount of data is increasing, they also become more structured. Metadata as well as structural relationships between results can give rise to a contextual overview of the data, allowing to filter them. Contextual overview can take many forms: topical, geospatial, and temporal being some of the more popular. Websites such as Flickr and Twitter provide rich data sources annotated with geospatial and temporal metadata. Users already familiar with searching in geographic environments like Google Maps, can find results presented directly on a map rather than in a list. There are also early prototypes for mapping the events and characters of cultural heritage onto maps and interactive timelines (e.g. Timeplot¹), driven by

enthusiasts using the Google Maps API² or the INSPIRE framework³.

These enthusiasts approached us to implement a web application that enables the synergetic exploration of multiple topical queries in a geospatial and temporal context. They want to be able to juxtapose and compare spatial distribution and temporal trends of multiple queries. We searched for other systems on which we could base our design. Among our major influences are a web visualization of *The Guardian* (Rogers, 2010), working with many datapoints, but supporting only one dataset and limited interaction, *GeoVISTA CrimeViz* (Roth et al., 2010), which shows three datasets at the highest zoom but aggregates them at lower zoom levels, and *VisGets* (Dörk et al., 2008), which supports single datasets with a wide variety of interactions. A general concern with these systems, which we selected as representatives of current systems for geospatial and temporal exploration, is also the clutter that arises when many data elements are spatially close together. This becomes especially harmful when the glyphs' attributes, e.g. which topical query the items arose from, differ. Also the time widgets in these systems can be improved as they typically do not support flexible queries like May 2009 to May 2010.

¹<http://simile.mit.edu>

²<http://code.google.com/intl/de/apis/maps/>

³<http://www.inspire-geoportal.eu>

In this paper we present the design of a novel web application that:

- combines showing and comparing multiple result sets in a geospatial and temporal context,
- improves geospatial presentation of items by zoom-dependent aggregation of items that avoids occlusions,
- allows more flexible selections within the map and time widget, and
- works client-side to allow user-provided data and improve response time.

2 RELATED WORK

There has been a lot of research on the usage of visual interfaces to interact with large data resources. Tukey illustrates the benefit of integrating interactive visual tools into the data exploration process to expedite the undirected search for structures and trends (Tukey, 1977). This so called Visual Data Exploration process follows the Information Seeking Mantra (Shneiderman, 1996): *Overview first, zoom and filter, details-on-demand*. This means, the first presentation of the information is an overview over the whole data, followed by the search for and drill down on interesting patterns, exploiting capabilities of the human visual system. The search for a specific information, which starts with a vague imagination, and in a stepwise process leads to understanding is called exploratory search (Marchionini, 2006).

Much research in the field of Visual Data Exploration was dedicated to the analysis of data with geographical and temporal metadata. An overview over tools, principles, challenges, and the concept of the analysis of geotemporal data are given by Andrienko and Andrienko (Andrienko and Andrienko, 2006; Andrienko and Andrienko, 2005).

Baldonado (Baldonado et al., 2000) proposed guidelines for the usage of multiple views in visualization to increase the user's ability to receive deeper insights if the data is shown under multiple aspects. Causal relationships can be found easier as well as unforeseen connections. Because of space restrictions we cannot address all but a few representative implementations that exploit synergetic effects from using linked views for geospatial-temporal data.

VisGets is a web application that displays the result of one query in four linked widgets (Dörk et al., 2008), working much like specialized dynamic queries (Shneiderman, 1994) for spatial and temporal filtering. A location widget shows the result items as small glyphs overlaid on a geographic map. A time

widget shows the distribution of results in bins representing years, months, and days. A tag widget shows a limited set of frequent words in alphabetic order and a size proportional to their importance. A results widget shows small textual or image thumbnails of results arranged as a table. *VisGets* offers query refinement in space (selecting one glyph), time (selecting a year, month or day bar), and by topic (selecting a tag). Each refinement affects the presentation in all other widgets. *VisGets* also provides many mechanisms to interact with the results of a single query, but a comparison of different queries is not possible.

GeoVISTA CrimeViz (Roth et al., 2010) for crime analysis enables comparison of different crime incident types. The map widget shows individual circles on higher zoom levels. However, on lower zoom levels, the incidents of all types are aggregated into hexagonal bins, thereby disabling comparison. The concept of binning items and plotting their total poses a problem, since the resulting continuous map overlay does not preserve the user's geospatial orientation. On the *CrimeViz* timeline, the number of incidents per time period of different types are stacked, which shows the sum of incidents and makes obvious differences between the incident types and global trends visible. Wang et al. (Wang et al., 2008) also uses a specific stacked graph – the *ThemeRiver* (Havre et al., 2000) – for the visual analysis of global terrorism activities. We avoid stacked graphs, because we do not want to sum up unrelated categorical data. Furthermore, it is hard to accurately interpret trends that lie atop other curves (Heer et al., 2010). *LISTA-Viz* (Hardisty and Klippel, 2010) uses the concept of small multiples (Tufte, 2001) for the analysis of spatio-temporal autocorrelations on maps and time-bars to allow comparison. The usage of small multiples makes comparison of our data, which typically has a highly irregular spatial distribution, difficult.

Maps can also benefit data of historic origin. Tshipidis et al. (Tshipidis et al., 2011) demonstrate the usefulness of interactive geospatial-temporal visualizations for the exploration of archaeological data. The *HESTIA* project (Barker et al., 2010) investigates the differences between imagined geographic distances and real distances in ancient Mediterranean space by mapping data onto contemporary maps: the *Herodotus Timemap*. It is a geo-temporal visualization that shows several location types (settlements, regions, physical features) mentioned in various text passages. The user is able to browse through time, space and book chapters. However, there is neither aggregation in space nor time, causing visual clutter. The system also lacks filtering of locations by dynamic refinement in space and time.

3 DESIGN

In this section, we describe the design of our system, starting from the input and the filtering methods to the representation inside multiple linked views and finish with the interaction techniques for query refinement. The presentation consists of a map widget showing the position of query results as small circles suitably aggregating items for the current zoom, a time widget showing the distribution of results in a time span, and a detail widget showing textual contents or thumbnails of data items arranged in a table.

The input to our system is the result of one to four result sets from a classical keyword or term-based topical search. Each item in the input sets is required to have a longitude and latitude as well as a time stamp or time span annotation. In our datasets the number of items that lack these data is sufficiently low so that we can simply omit them from the map and time widget.

The system provides spatial filtering based on containment queries in the form of circles and polygons as well as temporal filtering based on intersection queries with a given time span. Initially, the system starts with a minimally enclosing rectangle over all items' positions and a time span that stretches over all items' time stamps and time spans.

We choose colors to codify results from different datasets because of their effective use to discriminate categorical aspects of data (Bertin, 1983). We presume map colors to be mostly dark, cold, or unsaturated, and select very light colors for deselected and very saturated colors for selected circles to ensure that overlaid circles are not masked by map colors. We use a fixed number of colors, as the human ability to distinguish colors is limited. This also limits our system to four input datasets. From the 12-color palette suggested by Ware (Ware, 2004) for categorical usage, we chose to use only four colors because we wish to ensure that mixtures of these colors are distinguishable in our time-widget. The four colors used for datasets 1 through 4 are red, blue, green, and yellow. The order was chosen to reduce the chances of misinterpretation by color-blind people. We followed the recommendation of Harrower (Harrower and Brewer, 2003) for diverging colors in thematic maps. The user can also change the colors from their default values.

3.1 Map Widget

The dominant widget in our visualization in terms of screen space is the *map widget*. Following the definitions of Harris (Harris, 1999), the map widget consists of a geographical map, which is overlaid with bub-

ble glyphs. We use circles of varying size for glyphs, and each circle represents one or more items of the dataset. We chose a circle representation on the map to display data densities instead of, for instance, a heat map, for multiple reasons. Multiple result sets require multiple heat maps, which in turn would require mixing color maps, against which Ware argues (Ware, 2004). Also, humans are more accurate judging areas than they are judging color tones, making areas a better candidate for quantitative values (Cleveland and McGill, 1984). Using glyphs also allows to group glyphs and to make every data item individually accessible through its graphical representant.

In addition, we disallow the glyphs to overlap in order to avoid visual clutter. To achieve this goal, we merge circles based on their size, distances, and the current zoom level in an iterative process. We aggregate by proximity rather than geographic regions because regions can have a bad area-to-diameter ratio and the glyph resulting from the aggregation may extend beyond the border, leading to ambiguities.

Aggregation for One Dataset. In the following let N denote the number of items, k the number of supported zoom levels, and N_l the number of circles at zoom level l ($0 \leq l < k$). r_{min} is the radius used for circles representing one item. We define r_{min} dependent on the average font size of labels of common web mapping services (Google Maps, Bing Maps,...) so that small circles have salience equal to labels. The maximum radius is defined as $r_{max} = r_{min} + \log_2 N$. A_{min} and A_{max} are the resulting circle areas. The area A_i of an arbitrary circle c_i is interpolated linearly depending on the number n_i of data elements, that are assigned to it:

$$\phi_i = \left(3 - \frac{2l}{k}\right) \frac{n_i - 1}{N - 1}$$

$$A_i = \pi r_i^2 = (1 - \phi_i)A_{min} + \phi_i A_{max} \quad (1)$$

The weight in ϕ_i ensures larger area differences for higher zoom levels. The effect is, that differences between smaller circles are more salient.

To ensure non-overlapping circles, we perform a variant of the hierarchical agglomerative clustering algorithm. Two separate circles c_i and c_j need to be merged if their Euclidean distance d_{ij} after Mercator projection is smaller than the sum of their radii r_i and r_j . From this ratio we compute dissimilarities ψ_{ij} and merge circles c_i and c_j into a new circle c_{ij} placed at the barycenter of all their data elements if

$$\psi_{ij} = \frac{d_{ij}}{r_i + r_j + \varepsilon} < 1.$$

The gap value ε denotes the smallest permissible distance between two separate circle boundaries.




Figure 1: Aggregation of 307 points. The underlying *Guardian* dataset shows key incidents during the Afghanistan conflict. (a) Each circle with radius r_{min} represents one incident. The high degree of overplotting makes it hard to access each circle individually, and to determine the precise conflict centers. (b) 307 elements were aggregated to 43 non-overlapping circles. Larger circles clearly indicate conflict centers.

The aggregation of the circles starts with N circles from which a proximity matrix Ψ is generated. Initially, each single circle represents all items at a common point. Iteratively, we merge the two circles with the lowest value for ψ until there is no value below 1. Each merge step causes two deletions and one insertion of rows and columns in Ψ_{ij} . We end with N_{k-1} circles for the highest zoom level $k - 1$ and proceed similarly to compute the N_{l-1} circles at zoom level $l - 1$ from the circles at zoom level l for each $0 \leq l < k$. We perform this agglomerative hierarchical clustering procedure, which takes $O(N_{l+1}^2)$ time for each zoom level l (with $N_k = N$), until we have found a non-overlapping circle set for zoom level $l = 0$. An example result is given in Figure 1.

Aggregation for Multiple Datasets. In the case of multiple ($m = 2, 3, 4$) input sets, we compose multiple circles c_1, \dots, c_m into a more complex glyph – a *circle group* – whose bounding circle b will be used in the aggregation. We use a modified version of the congruent circle packing method of Kravitz (Kravitz, 1967): after computing the radius of each c_i with Equation 1 we place m template circles t_i at the same distance from a hypothetical center z that touch each other at their boundary. We then place the circles c_i inside the template circles t_i and move them as far as possible towards z without them leaving t_i . In the case $m = 4$ the biggest and second biggest circle are positioned opposite each other. We then construct a minimal bounding circle like in Figure 2.

Although the circle packing method requires a lot of screen space, it leaves all similarly located elements very close and quantitative differences that are reflected by varying diameters of the corresponding circles are easy to spot. Furthermore, all circles can




Figure 2: Example circle packings.




Figure 3: Level of detail of placename tag clouds. The used Flickr dataset shows photos which are tagged with “Italy”. (a) most of the places are “Italy” and “Vatican City”, (b) administrative divisions of Italy, (c) most photos for capital Rome, (d) Colosseum photos in Rome’s district San Paolo.

be accessed individually. We also chose packed circles over pie charts, because even though they can present ratios of the data at one point very good, it is difficult to compare the size of pie slices between pie charts of different size.

Placename Tag Cloud. Each of the map’s circles is optionally associated with a set of placenames including their frequency. A details-on-demand tag cloud for a circle shows the most frequent placenames, their font scaled proportional to their frequency. This feature provides a preview of how a glyph arising from agglomeration would split if zoomed in. If the data offers different levels of detail for a place, we choose the label dependent on the current zoom level. We distinguish four semantic levels: country, region (e.g., a state or a countryside), city, and borough (e.g, a district, specific place or an address of the given city). If one of levels is missing, we take the next coarser level if possible, otherwise the next finer level. An example of the placename tag cloud is given in Figure 3.

Historical Maps. Because some of our applicants wish to study datasets with historical context, we allow to show political borders for specific points of time. In addition to a contemporary map we offer 23 different historical maps from 2000 BC to 1994 AD provided by Thinkquest⁴. Because a dataset may span a time range for which multiple maps are available

⁴<http://library.thinkquest.org/C006628/>

and there is no concept of “average political border” we show the map closest to the median time stamp occurring in the dataset as default and allow the user to switch to any map at his convenience.

3.2 Time Widget

Using terminology of Harris (Harris, 1999), we construct the time widget as a segmented area graph, that has time as unit for the x -axis. The segmented area graph is a line chart where the area under the line is filled. T_1, \dots, T_n partition the interval $T = [t_{min}, t_{max}]$ of the given dataset into intervals of regular duration: either seconds, minutes, hours, days, weeks, months, quarters, years, or decades. The resolution is chosen to maximize the number of intervals without exceeding 400. Short units typically arise from steadily updated data sources and large units arise from data with historic content. The resolution unit changes automatically when the user zooms inside the time widget.

Whereas the x -axis of the time widget is directly defined by T_1, \dots, T_n , the y -direction shows the number of data items that fall in each interval using binning. For data items with time stamps the counting is straight-forward, for data items with time spans we add a value proportional to the amount of overlap with each bin. Although this can lead to an overrepresentation of items with long time spans in the time widget, we found this to be no problem. For the datasets we considered, either the effect was benign as the time spans had approximately the same duration or items with longer time spans were also more interesting.

Time Binning for Multiple Datasets. In the comparative setting we perform the bin counting per group. In the final visualization the bins’ sizes are shown as overlapping segmented area graphs rather than bar charts because the former is better suited to direct comparison of the groups’ time distribution. We use the same colors for the groups as in the map widget and shade the area under each line using a semi-transparent version of that color. This ensures that each curve is visible and also hints at the area that would have been present in the bar chart. Despite the use of blending, the limitation to four colors as well as the stacked area graphs’ shape helps limit and resolve ambiguities.

3.3 Detail Widget

For inspection of single data items that match the current filtering we present small textual or image thumbnails presented as a table in the detail widget. This

widget is the only one which does not include any aggregation, but results are presented on multiple pages if they exceed a certain fixed number.

3.4 User Interaction

Our application offers multiple ways for the user to interact with the data. Each widget provides native navigation that results in updates of the other widgets. The map and the time widgets provide simple zoom and pan. Because of the way we aggregate data in these widgets an animation between zoom levels or time resolution changes is not performed. When focusing or selecting items in one view the corresponding items are highlighted in the other widgets too.

The selections are linked. In particular, a selection within the map widget triggers a proportional selection on the time widget. This is done by overlaying the area graph with a bar chart that shows, how many elements of a timeslice are associated with the selected location. A selection on the timeplot triggers a proportional selection on the map as well. Therefore, we overlay a deselected circle with a circle with the highlight color and a radius, that reflects the proportion of elements, which are inside the selected time interval.

Highlight. A glyph or time step that the mouse hovers over is highlighted along in all other views. It is a preview for selection. Hovering over a circle on the map frames the circle with its associated placename tag cloud. For a hovered bin in the time widget, the number of associated elements and the corresponding date is shown.

Selection. Each highlight can be turned into a selection by a simple click. Every widget provides its own means of selection and offers a broad set of further user interactions, e.g., any selection can be turned into a query refinement by a context menu item.

In the map widget the user can select each circle individually. Thereby, the placename tag cloud is frozen and the different place labels can be highlighted as well as selected.

It is also possible to draw shapes such as circles or polygons. Selections can be refined by dragging shapes or control points. As a special selection the user may also click on a country effecting the selection of all the items inside its borders. In historical contexts, using maps of the time period of interest are of great benefit to this task.

The time widget allows both the clicking on one bin and the selection of a time range using a mouse drag gesture. A toolbar is then shown, that offers to

modify the left or right border of the selected time era and to add a “gray zone” which blends between selected and deselected elements. Finally, the user can manually move the selected time window. We update the map widget automatically to facilitate a direct reflection how locations change over time. A play button starts an animation mode and the selected time span loops smoothly over the entire time range. This mode directs the user’s attention towards changes inside the map.

The boxes in the detail widget reflect highlight and selection by setting the border and background, respectively, to the dataset’s color. Additionally, the table can be used for the selection and deselection of single items, to allow further refinement to the user.

Intra Dataset Comparison. One of the major features of the detail widget is the ability to export elements of a selection of a dataset as a new individual input set. Hence, the temporal comparison of different geographical regions of one dataset as well as the geographical comparison of different time periods is possible.


4 IMPLEMENTATION

We implemented our design in a prototype application in the context of the European project (Purday, 2009) *EuropeanaConnect*⁵. The system’s architecture is given in Figure 4(a).


The Web-server is mainly used for data processing. Dependent on the requested data source, the server retrieves the requested data and constructs a KML-file with a specific format. For each item that fits the given query, we add a *Placemark* tag to the KML root. Each Placemark is filled with a *name*, location in form of *latitude* and *longitude* and either a valid *TimeStamp* or *TimeSpan*. Optionally, a placemark node can be enriched with a location info (*address*), that will be used for the placename tag clouds. Slashes are used to separate their levels of detail. A *description* entry can be used to commit a CDDATA section that includes HTML content to offer detailed information for the placemark. Finally, the KML-file will be sent to the client as a response to the XML-HttpRequest.

The client browser parses the retrieved KML-file, creates JavaScript objects and fills the widgets with content. We decided to construct a pure JavaScript client, because it is supported by every browser without any additional package. Another reason for

⁵<http://www.europeanaconnect.eu/>



(a) Architecture.



(b) Screenshot.

Figure 4: System overview.


choosing JavaScript for our main application is the rapidly growing browser support in terms of performance. All interactions, including data refinement are performed on the client side. The great advantage of this system structure is, that each modification of a dataset within the client browser triggers only some functions on the client side. This benefits the response time of modifications, since the browser does not have to wait for the server. We allow the user to load his own KML-files. To the best of our knowledge our tool is the first to support this feature.

For the widgets on the client side we make use of two OpenSource JavaScript libraries. A modified *Simile Widgets Timeplot*⁶ instance organizes the segmented area graphs and *OpenLayers*⁷ visualizes the circles and arranges the different map layers. We offer some popular maps as well as the historic maps,


⁶<http://www.simile-widgets.org/timeplot/>

⁷<http://openlayers.org/>


COMPARATIVE VISUALIZATION OF GEOSPATIAL-TEMPORAL DATA




(a) CrimeViz (low zoom level): Incident types arson, homicide and sex abuse.




(b) Our system (low zoom level): Incident types arson (red), homicide (blue) and sex abuse (green) in circle groups.



(c) Selected incidents in 3rd police district.



(d) Compared crimes stolen car (red), burglary (blue), theft (green) and robbery (yellow).



(e) Co-located incidents at 27th of June.



(f) Stolen car and burglary incidents in neighborhood 13 (top) and 39 (bottom).

Figure 5: Crime incidents in Washington D.C. Figure 5(a) reproduced with permission from (Roth et al., 2010)

which are explained in Section 3.1. We use a web server instance of the *GeoServer*⁸, that provides the tiles of the historic maps. A screenshot of our application can be seen in Figure 4(b).

5 RESULTS

To demonstrate the flexibility of our system, we linked dynamic web sources as well as static data sets with our system and compared our visualization to existing ones. Finally, we performed a field experiment (McGrath, 1995) to evaluate our design.

5.1 Crime Incidents

We compared our visualization to the *CrimeViz* application. *CrimeViz* can only show incidents of one year at once; we chose 2009 for comparison. In *CrimeViz* the total of the frequencies of 3 incident types arson, homicide and sex abuse is displayed in form of hexagonal bins of different saturated red color hues on lower zoom levels (Figure 5(a)). Thereby, the overlay causes a loss of map context in strongly saturated areas. In our system places with many incidents are salient through circle groups with different sizes (Figure 5(b)). Without the need to zoom or filter, we can directly see that there was no homicide and arson incident in the north western neighborhoods. This cannot be seen in *CrimeViz*. Furthermore, the time resolution in our system is more precise (days) compared to *CrimeViz* (weeks). There are also no spatial filter mechanisms offered by *CrimeViz*. Within our system, we are able to filter incidents geographically by defining polygon or circle container shapes, or directly by selecting an administrative area, e.g., a police district (Figure 5(c)).

The used crime database⁹ offers a total of eight different crime incident types. *CrimeViz* used only the three types with smallest number of incidents: arson (59), homicide (140), sex abuse (361). Other types – theft (9,262), stolen car (4,852), robbery (4,389), burglary (3,670), assault with deadly weapon (2,623) – are not considered. Since each circle has one spatial representative in *CrimeViz* and there is no spatial aggregation on higher zoom levels, overplotting would be a fatal problem for the larger datasets.

The additional crime types allow us to compare crime incidents which may be more related to each other than arson, homicide, and sex abuse. In Figure 5(d) we see the geospatial and temporal comparison of stolen cars, burglaries, thefts, and robberies.

By hovering over the glyphs in the city center, we find out that there are relatively few stolen car (20%) and burglary incidents (23%) in comparison to thefts (45%) and robberies (35%). Moreover, we find patterns for thefts and robberies near populated places like metro stations or shopping centers. By exploring individual days, we discover that a stolen car incident is often grouped with a theft, burglary or robbery incident (Figure 5(e)). The number of correlations increases further by choosing a two day time range.

With the comparison feature, crime analysts could improve the detection of connected crime incidents of different crime types. Furthermore, our tool in combination with crime data could be used to drive the decision making for apartment search. Interesting questions could be “Where is a secure region in terms of burglary?” or “Where should I also rent a garage for my car?” Figure 5(f) indicates Neighborhood 13 as substantially safer than Neighborhood 39 with respect to stolen cars and burglaries.

5.2 Guardian Data

The *Guardian datablog*¹⁰ is another open source for a multitude of datasets with geospatial-temporal data. On 23rd of October 2010, *The Guardian* published the *Iraq war logs*. The data contains information about approximately 60,000 incidents during the Iraq conflict from 2004 to 2009. Every incident claimed at least one casualty; 109,000 total casualties are mapped. Each incident is annotated with a time stamp, latitude, longitude, incident type, and number of casualties for each casualty type. The *Guardian* also published a visualization (Rogers, 2010), where each incident is represented by one circle on a map (Figure 6(a)). The visualisation is cluttered as the result of overlapping glyphs. Figure 6(b) shows the geographical distribution of incidents with our system using the same color for glyphs. In both visualizations regions with a low glyph density work well, e.g., a series of horizontal circles from Baghdad (center) to Jordan (east) highlights the importance of this connecting road. However, the conflict centers, Baghdad, Al Mausi, and Al Basrah, cannot only be discriminated easily in our visualization.

The temporal distribution shows a peak of incidents between August 2006 and February 2007 (Figure 7(a)). The decreasing number of incidents afterwards indicates the success of the *Operation Imposing Law*, which started in February 14th 2007. The *Guardian* map shows information about the incidents' number of victims only in popups. We split the data into four sets, depending on the four given casualty

⁸<http://geoserver.org/display/GEOS/Welcome>

⁹<http://data.octo.dc.gov/>

¹⁰<http://www.guardian.co.uk/world/datablog>




Figure 6: Iraq War Logs: Incidents distribution. Figure 6(a) reproduced with permission from (Rogers, 2010). (a) Three conflict centers can be made out: Baghdad in the center, Al Mausi in the north, and Al Basrah in the south-east. Due to overplotting, it is difficult to determine which of these three regions suffers the most from incidences. The original visualization omits some incidences for unknown reasons. (b) Baghdad in the center is clearly the region with most incidences (23,783 incidences, approx. 5,600 incidences in Al Mausi, 2,400 in Al Basrah, as simple mouse-over reveals). All incidences are shown.

types (civilian, enemy, iraq forces, coalition forces) for comparison.

A comparison between incidents and civilian victims reveals two informations: (1) several incidents without civilian casualties in the north-west of Baghdad (Figure 7(b)), (2) incidents with many civilian casualties, e.g., an incident which caused 158 civilian casualties in the south of Karbala at April, 14th 2007 (Figure 7(c)). Another of these incidents is an IED Explosion in the north of the Baghdad's neighborhood Rasheed that caused 94 civilian casualties at February, 3rd 2007. The overplotting in the Guardian Map masks this incident: it is hidden by an event with one victim (Figure 7(e)). Our spatial aggregation procedure merges both incidents (Figure 7(d)); the incidents' details can be found in the detail widget.

The comparison of all casualty types shows that most of the victims were civilians (65,649), and the only period with more enemy than civilian casualties was at the beginning of the conflict in 2004 (Figure 7(f)). There was also less fighting at the road to Jordan in 2004. The roads to Mosul (north) and to Syria (north-west) were more contested at that time.


We are also able to reject the statement of the Iraq Military from March 14th 2007, that there were only 265 civilian casualties in Baghdad since the start of *Operation Imposing Law*. Figure 7(g) shows the time range from February 14th to March 13th with a total

of 997 civilian casualties and a peak on March 5th, caused by explosions around noon in Tigris.


5.3 Dynamic Data Sources

Humanities scholars within the *eAQUA* project (Büchler et al., 2008; Heyer et al., 2011) use our system to explore the geographic distribution and propagation of words extracted from texts in ancient Greek. The TLG research center at the University of California created the corpus, which is one of the most important resources that offers around 7200 Greek works of over 1800 Greek authors. Text mining algorithms extract words of strongly inflected ancient Greek and tag them with temporal and geospatial information. The metadata spans the eastern Mediterranean with Greece, Italy, Turkey, and North Egypt and 18 centuries starting from 8 BC. The resulting data divisions are shown in a suitable historic context by automatically displaying adequate historic maps. Figure 8 shows occurrences of the words "Plato" and "Aristotle" on the map of 400 AD.

We integrated our system into the online digital catalogue of Europeana. It contains over 15 million entries (images, texts, sounds, videos), which are digitized resources of Europe's libraries, museums, and archives. Most of the items have geospatial (e.g., publication place) and temporal (e.g., publication year)




(a) Peak of incidents 2006/07.




(b) Incidents (blue) and civilian (red) casualties.




(c) Incident near Kabala.




(d) IED explosion caused 94 casualties.



(e) Criminal event at same place (Guardian).



(f) More enemy (blue) than civilian (red) incidents 2004.



(g) 997 civilian casualties (top), explosions on March 5th 2007 (bottom).

Figure 7: Iraq war logs: analysis. Figure 7(e) reproduced with permission from (Rogers, 2010).




Figure 8: Occurrences of the words “Plato” (red) and “Aristotle” (blue) in ancient Greek texts: In the first period – Middle Platonism (1st to 3rd century) – we see a widely spread distribution of “Plato” and “Aristotle” (Figure 8(a)) in the Greek-Ionic region. The second period – Neoplatonism (4th to 6/7th century) – shows a lot of occurrences in the territories of Athens and Constantinople (Figure 8(b)). This indicates a correlated movement of both topics from rural regions to metropolises.

metadata. In the detail widget we use thumbnails, which link to the Europeana source page. Figure 9 shows the distribution of occurrences of different historic watermark types (tokens of letter authenticity).

5.4 Field Experiment

To evaluate our design and spot areas of usability improvements we conducted a field experiment. We asked 15 individuals to partake; 7 men and 3 women followed our call. The participant’s occupation ranged from pupil to pensioner and their age from 13 to 71 years. The participant’s educational background varied, but none had notable computer science background. Each test took place at the participant’s typical location where he accesses the internet. The average session duration was approximately two hours. A session consisted of three parts: a tutorial, solving tasks with the system, and an interview.

First, each participant was asked to work through a tutorial, which consists of 10 short flash videos linked from the application. The tutorial introduced the widgets and features of the system. To understand each feature, the users were encouraged to solve basic perceptual and manipulation tasks. Parts of the tutorial could be skipped or repeated, but we asked the participants not to ask questions to the present conductor of the study, who only observed, whether the pace, clarity, and volume of the tutorial was appropriate.

After finishing the tutorial, each participant was given a form with 15 prepared questions, which he had to answer using the system. During that time the participant could replay tutorials or ask the study conductor for help, but participants only used the opportunity to verify their understanding of the questions. The questions had an increasing complexity, and the

solutions required the use of all features mentioned in the tutorial. Questions could be answered in any order and skipped if considered to be too difficult. The solutions had different types: single or multiple data items, a specific time, place, placename, and even combinations of them. An example question asked to verify the Iraq police statement on the success of Operation Imposing Law (see Section 5.2). For some questions the users had to estimate solutions or find trends in space and time.

The fraction of correct answers was 75.3%. Even the oldest subject attempted 8 of 15 questions and all of his answers were correct. Some answers were incorrect but “almost” right (10.0%); there was often just a single overlooked glyph or time step. Only a small portion of answers were wrong (6.7%) or not solved (8.0%). By comparing the questionnaires, we discovered, that the ratio of correct answers was independent from the internet usage per day and the regularly visited platforms. Users with lesser experience required more time to finish the tutorial, and they also took more time in answering the questions, but the ratio of correct answers was comparable. Even experiences with web mapping services were not relevant for the correctness of the given answers.

Finally, the subjects were asked openly how they liked the application, followed by which features they liked or disliked, and what they would change. The general reception of the application was very positive. Although they stated to require the tutorial, which introduced only elementary interactions with single widgets, they learned and used the synergetic effects the more they used the tool. We also noted that the majority of the subjects did not make use of the time animation control, but those who took it to solve the questions, claimed it to be a useful feature.

Suggested improvements were mostly of technical nature, e.g., avoiding overlapping boxes inside the time widget and reducing the startup time of the system for very large datasets. Not all of them were using the browsers we recommended for fast JavaScript execution. They also proposed to use pie charts rather than circle packing, to provide placename tag clouds for circle groups, selections for continents, and a speed adjustment for the time animation feature. Some tasks, for instance the selection of multiple countries, have been implemented in the mean time.




Figure 9: Watermarks tagged with “Snake” (blue), “Bird” (red) and “Cloche” (green). Compared to Bird, Cloche and Snake watermarks were used more often. Whereas the Cloche watermark is spatially distinct and often used in the 14th and 15th century, the Snake watermark was used between 1450 and 1700, mostly in northern Italy and central Germany. Domain experts attribute this to the popularity of the Snake watermark as the heraldic device of the Milan dynasty of the Visconti, and later on as Swabian papermakers’ hallmark of excellence.

6 CONCLUSIONS AND FUTURE WORK

We presented a novel approach and web application to show, compare, and explore multiple topical query results in a geographical and temporal context. We were able to utilize, combine, and improve approaches from several related works. In comparison to CrimeViz (Roth et al., 2010), which also offers comparative visualization, we found a consistent solution for displaying results of different queries in the geospatial dimensions. We display them separately without aggregating them into the same representatives for coarser zoom levels. Furthermore, with our method we are able to solve the spatial overlap problem, which is also an issue in the Guardian visualization for the Iraq war logs (Rogers, 2010). Compared

to the similar system VisGets (Dörk et al., 2008), which only works for one set of results, we also made use of the linked widgets approach (map, timeline, detail widget) to extend the users exploration abilities. Furthermore, we enriched the filter capabilities in both geospatial (e.g., selecting all results inside a country) and temporal dimension (e.g., selecting dynamic time ranges). In contrast to VisGets, where most user interactions lead to a server request, we placed nearly all system logic inside the client to effectively improve the response times.

Our case studies show that visually comparable datasets extend the exploration and analysis abilities of the user in an effective way. It helps to detect equalities and varieties between distinct data contents that unveil their relationships in space and time. Our method is limited to four datasets at a time mainly to ensure that the colors used for discrimination are properly distinguishable, the splitting of circles does not waste too much screen space, and the overlapping segmented area graphs do not occlude each other too much.

In the field experiment, we could establish through the feedback of the participants, that our design is easy to learn and suitable to the given task of comparing multiple input sets in geospatial and temporal dimensions. In addition, the aggregation of items of distinct sets were evaluated as intuitive and the functionalities we offer to interact with the data were easy to adopt with the help of the provided tutorial – independent on the class of age, education level, and the amount of participants’ internet usage. We also received some hints on how to improve our application.

In the future we will direct our attention to the extension of our method to very large datasets by searching for client-server communication where most of the data remains on the server but a working set is transmitted to the client for quick interaction. We aim to extend our system to extract and show trajectories of distinct sets containing movement data, potentially leading to new insights.

The dynamic visualization of recent Twitter feeds is also one of our future objectives. We hope to see how topics proliferate. Preparing examples which automatically update their results by querying social platforms in equal time intervals could even show the live change of topics’ relations.

ACKNOWLEDGEMENTS

We like to thank our colleagues Christian Mahnke, Marco Büchler, and Muhammed Faisal Cheema for fruitful discussions and our field experiment partici-

pants for their time. We are also indebted to Cornelius Müller for preparing datasets. We also thank Simon Rogers and Robert E. Roth for providing screenshots of their systems. This research was funded by the European project EuropeanaConnect.

REFERENCES

- Andrienko, G. and Andrienko, N. (2006). Visual Data Exploration: Tools, Principles, and Problems. In *Classics from IJGIS: twenty years of the International journal of geographical information science and systems*, pages 475–479. CRC Press.
- Andrienko, N. and Andrienko, G. (2005). *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer.
- Baldonado, M. Q. W., Woodruff, A., Kuchinsky, A., and Kuchinsky, A. (2000). Guidelines for using multiple views in information visualization. In *Advanced Visual Interfaces*, pages 110–119.
- Barker, E., Bouzarovski, S., Pelling, C., and Isaksen, L. (2010). Mapping an ancient historian in a digital age: the Herodotus Encoded Space-Text-Image Archive (HESTIA). *Leeds International Classical Studies*, 9(1).
- Bertin, J. (1983). *Semiology of graphics*. University of Wisconsin Press.
- Büchler, M., Heyer, G., and Gründer, S. (2008). eAQUA - Bringing modern Text Mining approaches to two thousand years old ancient texts. In *Proceedings of the 4th International Conference on e-Science (IEEE08)*.
- Cleveland, W. S. and McGill, R. (1984). Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554.
- Dörk, M., Carpendale, S., Collins, C., and Williamson, C. (2008). VisGets: Coordinated Visualizations for Web-based Information Exploration and Discovery. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1205–1212.
- Hardisty, F. and Klippel, A. (2010). Analysing Spatio-Temporal Autocorrelation with LISTA-Viz. *Int. J. Geogr. Inf. Sci.*, 24:1515–1526.
- Harris, R. L. (1999). *Information Graphics: A Comprehensive Illustrated Reference*. Oxford University Press.
- Harrower, M. and Brewer, C. A. (2003). ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps. *The Cartographic Journal*, pages 27–37.
- Havre, S., Hetzler, B., and Nowell, L. (2000). ThemeRiver: Visualizing Theme Changes over Time. In *Proceedings of the IEEE Symposium on Information Visualization 2000, INFOVIS '00*, pages 115–, Washington, DC, USA. IEEE Computer Society.
- Heer, J., Bostock, M., and Ogievetsky, V. (2010). A Tour Through the Visualization Zoo. *Commun. ACM*, 53(6):59–67.
- Heyer, G., Büchler, M., Boehlke, V., and Schubert, C. (2011). Aspects of an Infrastructure for eHumanities. *ACM Journal on Computing and Cultural Heritage*.
- Kravitz, S. (1967). Packing Cylinders into Cylindrical Containers. *Mathematics Magazine*, 40(2):65–71.
- Marchionini, G. (2006). Exploratory Search: From Finding to Understanding. *Commun. ACM*, 49(4):41–46.
- McGrath, J. E. (1995). Methodology Matters: Doing Research in the Behavioural and Social Sciences. In *Readings in Human Computer Interaction: Towards the Year 2000*, pages 152–169. Morgan-Kaufmann, San Francisco, CA, 2nd edition.
- Purday, J. (2009). Think Culture: Europeana.eu from Concept to Construction. *Electronic Library*, 27(6):919–937.
- Rogers, S. (2010). The Guardian - Wikileaks Iraq war logs: every death mapped. <http://www.guardian.co.uk/world/datablog/interactive/2010/oct/23/wikileaks-iraq-deaths-map> (Retrieved 2011-10-07).
- Roth, R. E., Ross, K. S., Finch, B. G., Luo, W., and MacEachren, A. M. (2010). A User-Centered Approach for Designing and Developing Spatiotemporal Crime Analysis Tools. In Weibel, R. P. R., editor, *GI-Science*, Zurich, Switzerland.
- Shneiderman, B. (1994). Dynamic Queries for Visual Information Seeking. *IEEE Software*, 11(6):70–77.
- Shneiderman, B. (1996). The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Visual Languages, Proceedings*, pages 336–343.
- Tsipidis, S., Koussoulakou, A., and Kotsakis, K. (2011). Geovisualization and Archaeology: supporting Excavation Site Research. In Ruas, A., editor, *Advances in Cartography and GIScience. Volume 2*, volume 6 of *Lecture Notes in Geoinformation and Cartography*, pages 85–107. Springer Berlin Heidelberg.
- Tufte, E. R. (2001). *The Visual Display of Quantitative Information*. Graphics Press, second edition.
- Tukey, J. W. (1977). *Exploratory Data Analysis*. Addison-Wesley.
- Wang, X., Miller, E., Smarick, K., Ribarsky, W., and Chang, R. (2008). Investigative Visual Analysis of Global Terrorism. *Comput. Graph. Forum*, 27(3):919–926.
- Ware, C. (2004). *Information Visualization: Perception for Design*. Morgan Kaufmann.