Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets

Aidan Slingsby[1], Jason Dykes[1], Jo Wood[1], Keith Clarke[2]

1 – giCentre, School of Informatics, City University, London, EC1V 0HB, UK.
2 – Geography Department, University of California, Santa Barbara, CA 93106-4060, USA.

a.slingsby@city.ac.uk; jad7@soi.city.ac.uk; jwo@soi.city.ac.uk; kclarke@geog.ucsb.edu.

Abstract

'Tag clouds' and 'tag maps' are introduced to represent geographically referenced text. In combination, these aspatial and spatial views are used to explore a large structured spatio-temporal data set by providing overviews and filtering by text and geography. Prototypes are implemented using freely available technologies including Google Earth and Yahoo!'s Tag Map applet. The interactive tag map and tag cloud techniques and the rapid prototyping method used are informally evaluated through successes and limitations encountered. Preliminary evaluation suggests that the techniques may be useful for generating insights when visualizing large data sets containing geo-referenced text strings. The rapid prototyping approach enabled the technique to be developed and evaluated, leading to geovisualization through which a number of ideas were generated. Limitations of this approach are reflected upon. Tag placement, generalisation and prominence at different scales are issues which have come to light in this study that warrant further work.

1. Tags, Tag Clouds and Tag Maps

The exploration of large spatio-temporal data sets can benefit from the use of aspatial techniques developed for information visualization in a geographic context. Here we explore synergies between information visualization and geovisualization by using 'tag maps' [1] and tag clouds to help seek structure and relationships in the usage logs of a business directory for mobile telephones. We rapidly prototype these techniques using freely available technologies and reflect upon the possibilities for using this approach in the geovisualization of large datasets.

'Tags' are free form text labels that are independent of controlled vocabulary. They are widely employed for labelling digital content, such as photographs (Flickr[1]), video clips (YouTube[2]) and WWW bookmarks (del.icio.us[3]). Tags can form the basis of resource indexes, built by user communities for organising and sharing content due to their individual nature, widespread use and the diversity of those who contribute [2]. Patterns may emerge from these masses of tags, leading to classification schemes of tagged content known as folksonomies [3] that evolve over time in response to user communities [4].

Tag clouds are a visualization technique developed for assisting in this process by summarising the relative importance of tags [2]. Each tag is displayed, usually in alphabetical order, at a size according to some measure of its prominence. Tag clouds are now also widely used for summarising collections of words other than tags and online services exist for generating them[4]. Tag clouds may be extended to convey additional information about tags and the relationships between them. Hassan-Monteroa *et al.* [2] identify and order representative tags from tags clustered by semantic similarity, resulting in a tag cloud with a greater semantic range and in which the ordering is significant. Kerr [5] surrounds each tag with tags that share tag space, with a distance from the central tag reflecting the level of association. Techniques that rely on the distance and placement of words often use spatial metaphors for conveying relations [6]. Dubinko *et al.* [7] use the semi-spatial metaphors of the river and the waterfall to show temporal changes in the tag allocation of photographs in Flickr, and Havre *et al.* [8] use the river metaphor to show changes in the content of document

[1] http://www.flickr.com/
[2] http://www.youtube.com/
[3] http://del.icio.us/
[4] http://www.tagcrowd.com/

Slingsby, A., Dykes, J., Wood, J. and Clarke, K., "Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets", *Information Visualization 07, (E. Banissi, Eds.), pp. 497-504, IEEE Computer Society, Zurich, 2007.*

collections over time. Skupin [9] uses space more explicitly with 2D mappings of hierarchical similarities between documents.

Tag maps [1] are tag clouds in which the position of words is based upon real geographical space and sizes represent word prominences at specific locations – thus, tag maps can be considered to be tag clouds grounded in real geographical space (Figure 1).

Figure 1. Tag map (left) and tag cloud (right) of the top 20 business directory searches, centred on South Manhattan.

Milgram's "psychological map of Paris" [10] uses words corresponding to Parisian attractions and word sizes relating to the number of people who identified the attractions as such. This inspired Jaffe *et al.*'s [1] tagmap work, in which Flickr tags from georeferenced photographs are placed on a map. This is viewable through the Yahoo Tag Map applet [11] which enables users to explore the tags by zooming and panning the maps and to view some photographs from Flickr associated with the tags. Rather than Milgram's discrete items, these tags correspond to localised spatial concentrations of tagged items. Spatial distributions of this nature are often modelled as density estimation surfaces; e.g. Mehler *et al.* [12] show the spatial concentration of keywords used in news reports as estimated density surfaces shown as continuously shaded maps. Just as tag clouds have diversified to show the prominences of words other than tags, tag maps could show the spatial prominences of other words, e.g. Mehler *et al.*'s news report keywords.

The automatic placement of words on maps such that they are legible is a well-established cartographic research area [13]. Interactive maps provide the challenge of real-time placement of labels appropriate to users' interventions, but also the opportunity to dynamically reveal detailed information for user-identified neighbourhoods in real-time [14]. Such label-placement issues also apply to tag maps; in particular, how to deal with spatially-coincident prominent tags or words.

The prominence of tags for inclusion in a tag cloud is usually computed as the number of times the tag is used for tagging, but other measures are possible. It is often constrained by time (e.g. occurrences over the previous month). In the case of tag maps, prominence is necessarily constrained by space – i.e. the number of times a tag or word is used within a localised area – in the same way that spatial densities are a function of area [15]. Thus, prominence is explicitly linked to scale. In addition to their use for *visualizing* geo-referenced text, interactive tag maps and tag clouds can be used for *selecting* data in both geographical space and information (tag) space and exploring these scale effects (Figure 2) as part of the exploratory process. We can consider this as a cycle of filtering by space and attribute using the two linked views. An interactive tap map in which prominences are calculated in real time across a geography appropriate for the map's changing viewing extent helps explore scale effects. Complementing this functionality with a linked tag cloud constrained to the current viewing extent, reveals the extent to which concentrations of tags within the viewing are localised.

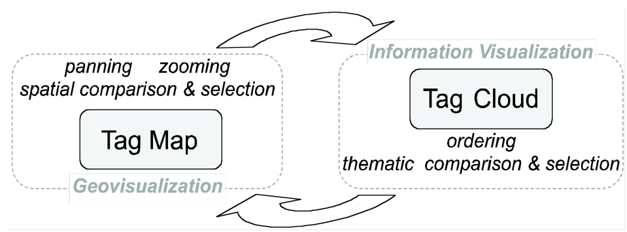

Figure 2. Exploratory cycle – Geovisualization proceeds using interactive views that are spatial and aspatial for overview and filtering.

These views and links are prototyped here using widely available technologies, which we evaluate informally in the context of exploratory geovisualization.

2. Data, Context and Approach

This work arose from a wider challenge to develop visualization techniques to explore a large data set with rich temporal, spatial and aspatial components. We also wanted to investigate the potential of combining a set of freely available network-friendly technologies that use *de facto* data standards with published APIs, for rapidly prototyping the techniques. This approach has become known as a 'mashup' [16]. Mashups often use mapping or graphical technologies as the basis for integration and can provide a more flexible alternative to single proprietary Geographic Information Systems for geovisualization. An advantage of this approach is that a prototype can be rapidly produced by selecting and flexibly combining a set of tools which work well together in a network-based environment. This approach has considerable potential for supporting the visualization process that requires flexibility and can benefit from rapid development and re-using data and code as the process progresses.

Slingsby, A., Dykes, J., Wood, J. and Clarke, K., "Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets", *Information Visualization 07, (E. Banissi, Eds.), pp. 497-504, IEEE Computer Society, Zurich, 2007.*

2.1 Data Set

A log of 1.42 million proximity-based requests for businesses and services made by mobile telephone users was supplied by go2[5], a US-based mobile telephone service. The requests were made over one month in 2005, and each request contained the business name queried and the locations and times at which each query was made.

2.2 Basic Requirements

The implementation of interactive tag maps linked to tag clouds requires the means to compute spatial and aspatial prominence with the ability to constrain by time, at a number of scales in real time, and at a speed that does not hinder visual thinking. Using these computed values, the means to position text at particular locations, at particular sizes and with some control over the colour symbolism used, is required. As the user zooms in and out, more words need to be revealed as the scale of spatial prominence changes. The linked tag map and tag cloud views required the means to trigger the generation of each view type in response to user actions.

2.3 Enabling Technologies

Two readily available spatial information clients were used to implement interactive tag maps – Yahoo!'s tag map applet[7] [1] and Google Earth[8]. The former is a Macromedia Flash applet, designed specifically for tag maps, that runs directly in most web browsers and which can read a stream of input data from a customised XML-based web service [11]. The latter is a widely-used general-purpose geobrowser that can display spatial and temporal data specified in KML [17] (an XML-based grammar). Geobrowsers such as Google Earth are helping enable the deployment of Internet GIServices [18] through their ease of use and access to high resolution geographic information [19].

Both browsers fulfil the basic graphical and user interface requirements stated in section 2.2. Both also provide contextual data and gazetteers; the former uses multiscale mapping and the latter provides access to a rich variety of additional high resolution datasets.

MySQL[9] is a free, fast and widely-used relational database, which together with the server-side scripting language PHP[10], fulfils the storage and computational requirements stated in section 2.2, enabling us to store 1.42 million records, retrieve spatio-temporal subsets and generate the XML-based outputs required by the Yahoo! Tag Map applet and Google Earth.

3. Implementation and Preliminary Findings

We present our technique and implementation and informally evaluate these through what we subjectively consider to be 'successes' assessed according to our expectations and basic requirements, before discussing limitations and associated issues that have arisen and which require further consideration. We also suggest some insights gained through the visualization. Typically, in the case of preliminary findings, insights relate to the nature of the data and the ways in which it models the phenomena of interest.

3.1 Successes

Storage and retrieval – The 1.42 million records were stored using MySQL, a task well within the capability of comparable relational databases. SQL provided the flexibility required to query, sort and retrieve data from these records and the speed to support real time interaction and visualization in both browsers.

Spatial aggregation – Jaffe *et al.* [1] calculated spatial prominence by using a spatial clustering algorithm for identifying local concentrations of tags. This technique involves starting with one tag, and iteratively grouping tags of the same value to form clusters until the maximum distance between the tags in the group is appropriate for a particular spatial scale. The iterative nature of this approach does not lend itself to its real time execution for any given spatial scale. Jaffe *et al* [1] therefore precomputed a hierarchical set of clusters to achieve the levels of interaction required.

In order to compute real-time spatial prominence, we implemented a computationally-efficient algorithm that counts tags of the same value contained within regular grid cells in the same way that a series of density estimation surfaces might be generated for each tag. This prominence measure is scale-dependent. The size of cells is arbitrary, but computing the cell width as a proportion of the width of the viewable map display allows the scale dependence to be interactively explored. The algorithm was implemented as an optimised SQL query that processed only those records that lie within the currently viewable geographical extent.

5 `http://www.go2.com/`

7 `http://developer.yahoo.com/yrb/tagmaps/`

8 `http://earth.google.com/`

9 `http://www.mysql.com/`

10 `http://php.net/`

Speed – MySQL and PHP enabled us to dynamically generate tag maps and tag clouds at any spatial scale in response to user interaction. Both browsers can report their viewable geographical extents, which are used by PHP to set the cell resolution and query the database. The database ran at 2.4GHz on a Pentium Xeon dual core, and after optimising the queries and building an index, a tag map could be generated in less than one second for high zoom levels and in ten seconds for the entire country – speeds that we found acceptable for the interactive exploration that is fundamental to visualization.

Cartography / Representation – Both clients enabled us to place suitably sized text at specific geographic locations and to update this information as the maps were panned and zoomed (Figure 3). The XML language used by Yahoo!'s tag map applet provided just enough control over text formatting for our requirements.

Figure 3. Tag maps of the most prominent word in each (identically-sized) grid square; Yahoo!'s tag map applet (left) and Google Earth (right).

KML [17] allows much more control over formatting, including a wider range of text sizes, colour and transparency. We reduced the opacity of the more prominent (larger) words in the tag maps – a useful technique for developing visual hierarchies in dense and overlapping text. Semi-opaque backdrops were used to diminish the prominence of the aerial photography and ancillary spatial data and to visually synthesize tag maps with continuous data such as the density surface of tag prominence (Figure 4).

Figure 4. Tag map of the 20 most prominent business requests and their aggregate densities as a high resolution raster surface.

Ancillary data – Both browsers have contextual data and built-in gazetteers. The Yahoo! Tag Map applet has street-level mapping and Google Earth has access to a range range of data including high resolution aerial photography, administrative boundaries, geo-referenced photographs[11] and elevation data. This ancillary data, considered through visual synthesis, contributed significantly to our understanding when visually analysing the go2 data set and resulted in the development of a number of ideas.

Linking and interactions – Panning, zooming and interacting with the various graphics is accompanied by automatic map updates in both viewers. This supports the visual exploration of tags in their geographical context (tag map), and provides an information summary of selected geographic areas (tag cloud). This is where the limits of interaction in Yahoo's tag map applet are currently reached. In Google Earth, hyperlinks in tag maps can be used to trigger the generation of HTML tag clouds that are displayed in Google Earth's integrated web browser. Hyperlinks in the tag clouds can also trigger further HTML or KML. It is this hyperlinking between tag maps and tag clouds that facilitates the cyclical exploration process shown in Figure 2.

Figure 5. Filtering by geography and attribute.

For example, Figure 5 shows the map and spatial tag cloud for San Francisco. Larger words in the tag cloud

[11] http://www.panoramio.com

Slingsby, A., Dykes, J., Wood, J. and Clarke, K., "Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets", *Information Visualization 07, (E. Banissi, Eds.), pp. 497-504, IEEE Computer Society, Zurich, 2007.*

draw attention to potentially interesting patterns or outliers. Here the details of the '*california check cashing store*' tag can be inspected by clicking on the tag in the tag cloud, triggering the generation of KML that zooms to the spatial extent of this tag in the tag map (lower image), a small area in this case. The tag map and tag cloud can be subsequently updated. Inspecting the temporal properties of this tag through additional graphics showed that these requests were made at exactly the same time from the same location. This leads to questions about information seeking behaviour, its representation in the database and ways of associating records with the needs and behaviours of individual users [20]. The updated tag map and tag cloud can be further queried to promote additional insights into the nature of the spatial distribution of tags – particularly important when exploring large datasets that require aggregation, selection or summary.

Timestamps can be associated with KML elements, causing Google Earth to display a scaled timeline that can be used to interactively filter data by date and time (see upper right of images in Figure 6). For example, tags of log requests can be constrained by the time at which they were made using the timeline. It is also possible to map other ordinal data to the timeline. Figure 6 maps the days of the week as dates, but original times are retained. Sunday is 10th January through to Saturday as 16th January. The tag map for Friday evening to early Saturday morning (18:01 to 06:23, top) shows service requests that focus on eating and accommodation whilst on Saturday between 09:22 and 14:03 (bottom) the emphasis is on finance and jobs at home.

Figure 6. Interactive timelines for exploration. Tags are constrained to Friday night (top) and Saturday morning and early afternoon (bottom).

Figure 6 also illustrates how further symbolism can be introduced to convey additional information about tags in tag maps. Upper case is used to show positive deviations from expected occurrences of words based upon their spatial and temporal aggregate ratios. This and other information can also be conveyed in colour (specified in KML), using a continuous rather than binary encoding.

3.2 Limitations / Issues

The mashup approach has enabled us to rapidly prototype a set of techniques, to test different solutions and to identify issues that warrant further research. Although we have found that Yahoo!'s Tag Map applet and Google Earth provide accessible, useful, and flexible functionality, we have also found some limitations. In particular, the balance of map legibility with positional accuracy and data omission is an issue that requires further consideration.

Text placement – The grid-based approach computes the prominence of words for each cell as simple frequencies. As many words as can be displayed legibly should be shown in the resultant tag maps. Yahoo!'s Tag Map applet displays all the words it is given, often resulting in completely illegible maps (Figure 7, top left). This problem was addressed in Figure 3 by only displaying the most prominent word in each grid cell for display, but as the figure shows, this does not necessarily prevent the map from being cluttered because it is hard to predict whether words will overlap. Google Earth reduces cluttering by automatically and selectively culling words. Those with the same formatting style (size, colour, opacity) are not allowed to overlap.

Figure 7 (top left) shows that plotting words at the centres of grid squares results in extensive illegible overplotting in Yahoo!'s tag map applet and mass culling in Google Earth (top right). Randomly locating words within cells gives much better results in Google Earth as more space is available and so fewer words are culled (centre right). A Gaussian distribution around grid centres was also tested to address possible edge effects, which may help interpretation by reducing the likelihood of words being placed close to cell boundaries (bottom). However, the spatial focus of the Gaussian distribution may result in more words being culled by the Google Earth browser than occurs with the more dispersed rectangular random distribution. There is evidently a cartographic trade-off here that can be overcome to an extent through interaction and by providing alternative versions of the tag maps. A grid showing the boundaries of cells used in the aggregation, provides important information about word distribution (implemented in Google Earth as a KML 'layer' – see Figure 7, right).

Slingsby, A., Dykes, J., Wood, J. and Clarke, K., "Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets", *Information Visualization 07, (E. Banissi, Eds.), pp. 497-504, IEEE Computer Society, Zurich, 2007.*

Figure 7. Word placement in Yahoo!'s tag maps applet (left) and Google Earth (right). Words are located at cell centres (top), randomly within cells (centre) and using a Gaussian random function around cell centres (bottom).

Culling – Overplotting can be addressed by modifying the position of tags or by culling tags when generating KML. We have also seen that Google Earth employs a tag culling algorithm to reduce clutter. In this section, we refer to the effects of this when the maps is zoomed and panned using the *same data*, i.e. in a mode in which the tag map is not automatically updated for the new viewing extent. The way in which Google Earth selects words for culling is not documented and we found a lack of consistency in tags displayed chosen between zoom levels and panned views. Figure 8 shows the same KML file at different Google Earth zoom levels, in which the sample of visible words is very different, though this discrepancy is less marked where fewer words compete for space. A similar effect is evident when panning.

Dealing with culled and obscured tags – The Yahoo! Tag map applet displays all the tags it is given, however illegibly. This might be used as an indication that there is too much information to synthesise all at once and that one needs to zoom or filter the results. However, the issue is left to the XML web service; e.g. the web service used in Figure 3 (left) only supplied one tag per grid square. Google Earth's tag culling algorithm operates dynamically as it places words on the map. As discussed, this is achieved using means that are undocumented without indication to the user and produces results that are not consistent between views. This is of particular concern if the culled selection is as a result of some artefact of the data such as the word length or the first letters of words, because this would introduce undesired bias. If culling were completely random then every static view of a tag map might be seen as a random sample. Were this the case, repeated panning, zooming and tag map regeneration would result in the user being exposed to multiple random samples, revealing the degree of culling taking place and the diversity of tags.

Figure 8. The same data at different zoom levels. The lower view is zoomed to the box extent.

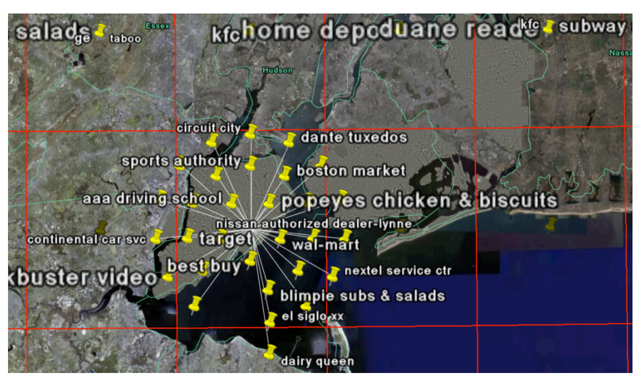

Figure 9. The user-initiated inspection of all tags in a specific grid square.

Slingsby, A., Dykes, J., Wood, J. and Clarke, K., "Interactive Tag Maps and Tag Clouds for the Multiscale Exploration of Large Spatio-temporal Datasets", *Information Visualization 07, (E. Banissi, Eds.), pp. 497-504, IEEE Computer Society, Zurich, 2007.*

Alternatively details can be provided on demand. Google Earth can automatically collapse multiple spatially coincident items, 'exploding' these when clicked [14]. In Figure 9, grid-squares have symbols at their centres. When these are clicked all the tags contained in the grid square are revealed.

Culling and spatial dilution – Figure 10 shows a tag map and a linked tag cloud of the twenty most popular text strings occurring in queries centred on South Manhattan. The tag map is dominated by a large '*duane reade*' tag, but '*starbucks*' is not visible, despite being prominent in the associated tag cloud. Either the '*starbucks*' tag has been culled by the browser or it may be the result of its spatial dilution (low frequency over a wide area). Generating a tag map of these two key words shows that both explanations play a role, but the main effect seems to be that the sources of '*starbucks*' queries are more widely distributed, suggesting that our tag map and tag cloud combination has identified a geographic effect. Using both the tag map and tag cloud in combination enables us to explore the distributions of these words in geographical space and information space and identify differences in the geographic nature of particular queries.

Figure 10. Spatial dilution. Tag map (left) and associated tag cloud (centre) with tag map of only '*starbucks*' and *'duane reade'* tags (right).

4. Conclusions and Ongoing Work

We enhance tag clouds with tag maps and use freely available network-friendly technologies to dynamically link the two for selection and filtering. By spatially 'grounding' textual information, we show how data can be filtered and selected in both geographical space and in tag space to reveal characteristics of each at different scales and at different times. Tag maps allow us to browse and generate spatial tag clouds. Tag clouds allow us to select aspects of information space and see the geography. Technologies such as MySQL, PHP and spatial XML dialects provide simple yet powerful means of doing so. Yahoo! Maps and Google Earth are examples of freely available spatial browsers that make a wide range of geographical information available for visual synthesis to empower and stimulate geovisualization. Such tools are updated regularly and possibilities change rapidly resulting in unexpected opportunities, but also some risks when undertaking geovisualization in this way. For example, at any time, more control may be provided over name placement priority in Google Earth and Yahoo! may provide additional options for text symbolism. Both organisations are associated with active user communities with which they appear to be well engaged. Consequently, specific results reported in papers such as this are subject to change (and may indeed affect change). Our key findings, however, relate to the more general opportunities for synergy between information and geovisualization to support exploration and the development of bespoke visualization techniques such as interactive tag maps and tag clouds using widely available technologies. The configuration we describe has helped us identify structure and anomalies, detect spatial and temporal patterns and effects and to develop hypotheses in our exploration of a large multivariate spatial dataset in which text plays an important role. The ancillary data available in Google Earth was particularly useful for context and ideation prompted by visual synthesis.

Work continues on techniques that support geovisualization using this technological configuration. The cell-based approach for generating tag maps is computationally efficient and can be computed for any required spatial scale – important characteristics for real time multi-scale visualization. However alternative non-raster approaches to spatial aggregation may address some of the artefacts of grid-based sampling. Hierarchical schemes, such as the Hungarian clustering algorithm mentioned in section 1.3, are potentially useful sources of information about relationships between spatial scales – investigating cluster composition, for example. In terms of cartography, colour can be used to show characteristics of words other than their under or over-representation in a spatial sample. For example, symbolizing temporal frequency or geographic extent can be achieved and may be useful. Whilst we have developed methods to address some of the browser behaviours in terms of the selection and culling applied, our investigations into these are ongoing and involve isolating the effects and assessing their implications. One approach is to let the user control spatial tag sampling where tags overlap and employ a multiple random method where tags on a tag map are particularly dense. Our tag clouds are arranged alphabetically, but other orderings may be beneficial [2,7]. Tag maps and tag clouds can be used to explore a range of text-rich spatial data sets including spatially referenced news items [12], the origins of family names[12] and spatial tags [1]. KML can be used to generate abstract

[12] http://spatial-literacy.org/UCLnames/

spatial graphics representing the multivariate nature of numeric datasets as well as textual information for display in Google Maps and Google Earth.

Importantly, the open and accessible methods reported here provide the flexibility to develop such bespoke combinations of layout, symbolism and interaction for visualization in a way that cannot be readily achieved with off-the shelf GIS software – interactive tag maps were rapidly and effectively implemented for real time geovisualization. The specific techniques presented here allow us to visually explore the relationships between the frequencies, times and locations of the requests, and through these, to interactively constrain the data by spatial and temporal extent for visual analysis. The interactive tag maps allow us to explore patterns at different spatial and temporal scales. Linking the tag clouds to the tag maps provides a non-spatial and non-temporal view of the same data. Our preliminary findings indicate that this combination provides an intuitive and powerful means to explore a large spatio-temporal dataset with a strong textual component with methods drawn from information visualization and geovisualization.

Acknowledgements

We gratefully acknowledge the support of Lee Hancock and go2 Directory Systems (http://www.go2.com) for permission to use the query data and help with this work. The Leverhulme trust funded Keith Clarke's sabbatical at City University. Four anonymous reviewers provided informative suggestions that have resulted in significant enhancements to the submitted draft of this paper.

Large, high resolution colour figures

`http://www.gicentre.org/papers/gIV07/`

References

[1] Jaffe, A., and Naaman, M., "Generating Summaries and Visualization for Large Collections of Geo-Referenced Photographs", *MIR 2006: 8th ACM SIGMM International Workshop on Multimedia Information Retrieval*, Santa Barbara, CA, ACM, 2006.

[2] Hassan-Montero, Y., and Herrero-Solana, V., "Improving Tag-Clouds as Visual Information Retrieval Interfaces", *InScit2006: International Conference on Multidisciplinary Information Sciences and Technologies*, Mérida, Spain. 2006.

[3] Hammond, T., Hannay, T., Lund, B., and Scott, J., "Social Bookmarking Tools (I)", D-Lib magazine, v11(4). 2005.

[4] Russell, T., 2006, "Cloudalicious: folksonomy over time", *6th ACM/IEEE-CS Joint Conference on Digital libraries (JCDL '06)*, Chapel Hill, North Carolina, p. 364, 2006.

[5] Kerr, B., 2006, "Sketches: worth a thousand words: TagOrbitals: a tag index visualization", *ACM SIGGRAPH 2006 Sketches SIGGRAPH '06,* Boston, Massachusetts. 2006.

[6] Wise, J.A., Thomas, J.J., Pennock, K., Lantrip, D., Pottier, M., Schur, A., and Crow, V., "Visualizing the Non-Visual: Spatial Analysis and Interaction with Information from Text Documents", *in Card, S.K., Mackinlay, J.D., and Shneiderman, B., eds., Readings in Information Visualization: Using Vision to Think*, San Francisco, Morgan Kaufmann, pp. 442-445, 1999.

[7] Dubinko, M., Kumar, R., Magnani, J., Novak, J., Raghavan, P., and Tomkins, A., "Visualizing Tags over Time", *WWW 2006*, Edinburgh. 2006.

[8] Havre, S., Hetzler, E., Whitney, P., and Nowell, L., 2002, "ThemeRiver: Visualizing thematic changes in large document collections". *IEEE Transactions on Visualization and Computer Graphics,* v8, p. 9-20. 2002.

[9] Skupin, A., 2000, "From Metaphor to Method: Cartographic Perspectives on Information Visualization", *IEEE Symposium on Information Vizualization,* Salt Lake City, Utah, p. 91. 2000.

[10] Milgram, S., 1977, "Psychological Maps of Paris, The Individual in a Social World", Addison-Wesley, pp. 68-90.

[11] Yahoo! Inc., "TagMaps Web Services : Roll Your Own API – Yahoo! Developer Network". http://developer.yahoo .com/yrb/tagmaps/tagapi.html

[12] Mehler, A., Bao, Y., Li, X., Wang, Y., and Skiena, S., "Spatial Analysis of News Sources", *IEEE Transactions on Visualization and Computer Graphics*, v12, pp. 765-772. 2006.

[13] Van Dijk, S., Van Kreveld, M., Strijk, T., and Wolff, A., "Towards an evaluation of quality for names placement methods", *International Journal of Geographical Information Science*, v16, p. 641-662. 2002.

[14] Fekete, J. D. and C. Plaisant. "Excentric Labeling: Dynamic Neighborhood Labeling for Data Visualization". *CHI 99: the CHI is the limit, human factors in computing systems*, Pittsburgh; PA, New York. 1999.

[15] Goodchild, M., and Quattrochi, D., "Scale, multiscaling, remote sensing and GIS", *in D.A. Quattrochi and M.F. Goodchild (eds), Scale in Remote Sensing and GIS*, London, CRC Press, p. 1-11. 1997.

[16] Miller, C. "A Beast in the Field: The Google Maps Mashup as GIS/2", Cartographic, v41(3), pp. 187-199, 2006.

[17] Google Inc., 2007, KML 2.1 Reference – Google Earth. http://earth.google.com/kml/kml_tags_21.html

[18] Zhang, T., Clarke, K.C., and Guan, Q., "The GIS Web Portal: Beyond Data Services", *AutoCarto,* Vancouver, WA. 2006.

[19] Clarke, K.C., Guan, Q., and Zhang, T., "The Anatomy and Physiology of the GeoGrid", *AutoCarto*, Vancouver, WA, 2006.

[20] He, D., Goker, A., Harper, D., "Combining evidence for automatic Web session identification", *Information Processing and Management 38(5)*, pp. 727-742, 2002.