Visualising Context and Hierarchy in Social Media

Suvodeep Mazumdar, Fabio Ciravegna, Anna Lisa Gentile and Vitaveska Lanfranchi

OAK Group, Department of Computer Science, University of Sheffield, 211 Portobello, Sheffield, United Kingdom
{s.mazumdar,f.ciravegna,a.l.gentile, v.lanfranchi}@sheffield.ac.uk
http://oak.dcs.shef.ac.uk

Abstract. The amount of data available online for analysis and reuse has dramatically increased in the last few years, with the advent of Linked Data and the mass adoption of social media. This has created a new class of users that want and need to consume available datasets to support their information needs but do not possess the technical expertise or the time to use complex querying systems. An example is Emergency Responders (ERs), teams of specialised users that monitor, react and respond to emergencies. Whilst ERs want and need to use these new sources of information, they have very stringent requirements on the type of interface that suits their information needs: the interface needs to be intuitive, highlight important data at a glance, and allow them to understand the context and the specificity of the data. Our contribution to this issue is a new visual analytics approach that focuses on context and specificity of social media content. We extract contextual and hierarchical features from social media messages and use them to provide co-occurring and hierarchical structures to explore large scale, unknown datasets in a highly visual and interactive manner. We introduce a new visualisation paradigm, called Context and Hierarchy chain, that provides means to explore such multidimensional datasets.

Keywords: visualization, context, hierarchy, social

1 Introduction

The vast amount of data currently available from Social Media (SM, i.e. Twitter, Flickr etc.) and Linked Data (LD) is a mine of potentially invaluable information in many domains: medical science, education, public administration etc. This has paved the way for a new class of users that have very specific information needs that are often reflected in the available data but hard to satisfy. Such users are often highly knowledgeable of the domain but lack the technical skills to query the data with existing interfaces. Available tools to browse semantic data are often too complex, requiring the user to form semantic queries, and often text-based [7]. On the other hand, existing interfaces presenting social media data are easier to use but often lack the added value of semantics: users are faced with a constant stream of new information and very few mechanisms to understand it.

Our approach targets the Emergency Response domain as one of the application areas[1] since this is a domain where the timely availability of up-to-date information is key to achieving a task. SM data and LD have become a natural port of call to achieve Situational Awareness (SA, i.e. accurate, complete and real-time information about an event) during an emergency or an event. They make available up-to-date, large scale amounts of data that can help ERs assess a situation. Information in social media is timely: Sakaki [19] noted as the first tweets to warn of an earthquake happens within a minute from the event. Information is large-scale: during the Superbowl 2011 final around 12.2 million related tweets were generated, an average of more than 20,000 a second[2]. Information is also distributed across several users. However, the data is often noisy and of unknown content.

The goal of our research is to boost the users' understanding of an event by the timely visual analysis of large-scale social streams (aggregated from Twitter, Facebook and Flickr) and linked data[3]. To this end, we have designed and developed a new visualisation approach that leverages the semantic value of the information, to create contextual and hierarchical structures for visually browsing information. The aim is to make the datasets more accessible and intuitive by presenting different levels of information (hierarchy) and correlated information (context) in a unique display. Our approach is inspired by the Level of Detail (LoD) methodology [9], where the users can switch between overview and detailed views but adds semantic features on which to cluster, zoom and filter, namely context and hierarchy.

In the following sections we discuss design recommendations for ERs and how these have been translated in a visual exploration system. The methodology for obtaining and visualising the data is described, with a focus on the new visualisation widget we designed. Details about the implementation and a comparison with the State of the Art are then presented, followed by Conclusions and Future Work.

2 Design Recommendations

For the purpose of our research we used the definition of Levels of SA provided by Endsley [8] and we analysed the requirements for each of the levels, translating them in design recommendations for a visual interface. The design recommendations were elaborated by analysing the three Levels of SA with experienced UK Emergency Responders during a series of individual interviews, focus groups

[1] Several other application areas and use cases have also been identified such as news broadcasting, public health, disaster management, politics and so on.

[2] http://mashable.com/2012/02/06/super-bowl-xlvi-social-tv-stats/

[3] Our exploitation of linked data is presently limited to understanding the hierarchical structures of the concepts that the posts are related to

and workshops conducted as part of a large European Project (WeKnowIt[4]) and a UK funded project (TRIDS[5]):

- **Level 1 SA** The first level of Situational Awareness involves ERs perceiving the different elements in the environment. The user requires to clearly understand the status of the most relevant elements in the context. These requirements are translated into design recommendations such as using familiar visual metaphors, providing filtering mechanisms to enable focus on an element of interest, highlight trends and spikes in the data and provide simple separate displays to focus on an element.
- **Level 2 SA** The second level of Situational Awareness requires ERs to comprehend the present situation. At this level, the users have recognised the elements that are important in the context of their predefined goals. Design recommendations for this level are presenting up to date information in the right spatial and temporal contexts, providing multiple coordinated displays that facilitate analysis.
- **Level 3 SA** The final level of Situational Awareness is the most challenging one, involving projection of future status and anticipating the consequences of elements on the context. The design recommendations that are translated from these requirements are providing mechanisms to highlight and follow up correlations between multiple elements at different levels of granularity, clustering elements to highlight implicit relations, integrating multiple features into a single display, providing multiple access and exploration points and providing flexible pathways for exploring related information.

3 Methodology

The design recommendations presented in the previous section have been achieved using a two-step methodology:

- **Data Processing**: this step analyses SM data to semantically enrich them and extract the key features to satisfy the information needs and the design recommendations.
- **Data Visualisation**: this step takes as input the semantically enriched SM data and applies the LoD paradigm to multiple dimensions to create a visual analysis framework that meets the design recommendations.

In the following sections we will describe the two steps, focusing our attention on the Data Visualisation step, as this is the core subject of this paper.

3.1 Data Processing

SM messages are typically composed of metadata (e.g. about the users, the equipment used to post, the location, etc.) and content, typically in textual format.

[4] http://www.weknowit.eu/

[5] The TRIDS project was a WeKnowIt follow-up project, under the UK Technology and Strategy SBRI programme call ‘Have I got ‘Views’ for you?’

Both metadata and content can be analysed to extract information (e.g. keywords, terms, named entities, events, etc.) and to create semantic data. In our approach, the messages are processed using Natural Language Processing (NLP) techniques which utilise the historic and current data and external knowledge resources (e.g. Wikipedia, OpenStreetMap) to augment the messages with structured metadata. Each message is annotated with tags and tags are resolved to Wikipedia concepts. This creates additional dimensions in the data:

- extracted tags can be used to establish the context of the message
- semantic concepts have an intrinsic hierarchy that allows exploration at different levels of detail

The following example shows how a tweet is enriched following a data processing stage. The user, date, geolocation and so on are extracted from the metadata of the tweet, whereas the content is analysed to identify ontological concepts (here, Sheffield being an instance of the city concept within the DBpedia[6] ontology).

```
<User>The Star, Sheffield</User>
<Date>20/09/2012</Date>
<ontology:city>Sheffield</ontology:city>
<Tweet>
Give your backing to Sheffield venues in running for top awards:
Tramlines is encouraging everyone to get behind... http://bit.ly/VfBrM4
</Tweet>
```

3.2 Data Visualisation

Our approach is inspired by the LoD methodology, where the aim is to "provide a balance of local detail and global context" [9]. We made use of the LoD principles to implement the design reccomendations. Level 1 recommendations were achieved by creating high level overviews using contextual features extracted from the Data Processing step such as Authors, Cited users, Date, Location, Source, Language, Tags, Sentiment.

Using the extracted features we created separate displays and we added filtering mechanisms using a standard faceted browsing interaction paradigm[10, 21]. These high level views provide easy ways to focus on an element of interest and gain an understanding of the most relevant elements. For example, a tag-cloud based visualisation (bottom right, Fig. 2) is used to highlight trending topics or users. Several data features are used to calculate trends and display them in multiple widgets: for example for each topic the most active and most mentioned users are calculated and displayed. Figure 1 (a) shows an example of time-based visualisation that highlights the frequency distribution of a tag to reveal possible spikes.

Level 2 recommendations were achieved by contextualising the information. For example, a spatial density visualisation was created to understand which are the most relevant locations for a topic or a user. This has been created by

[6] http://dbpedia.org/About

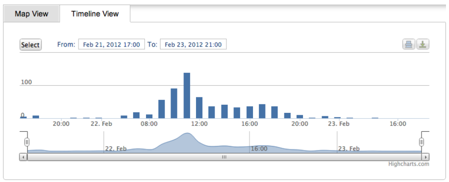

(a) Clustering of social media posts geolocated around london for the tag "999exercise" (a security simulation exercise in preparation of London Olympics). The time series shows exactly when a certain tag emerge and when it spikes, facilitating the exploration of relevant timeframes

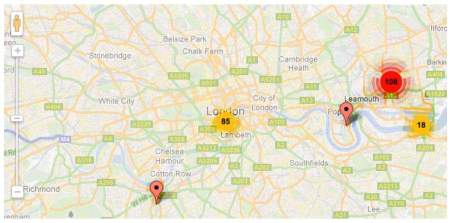

(b) Frequency distribution of the tag "999exercise" . The map visualisations empower users to rapidly assess the geographical spread of information and to focus on specific "hot" areas.

Fig. 1: Geographical map and timeline view provide ways for users to explore spatial and temporal dimensions of social media content. The data used here is an aggregation of Twitter, Facebook and Flickr posts, during the London 2012 Olympics security exercise (http://www.bbc.co.uk/news/uk-17817370)

plotting on a map the extracted locations and using a clustering algorithm for aggregating data and displaying frequency counts. Multiple visual parameters are used to map different attributes to the visual display: e.g. different colours are associated with the density of information, conveying how different regions are affected/involved in an event (see Figure 1 (b)).

Users can then delve into details by manipulating each visualisation widget to drill down and analyse the data accordingly to specific foci of interest. A time series visualisation is used to visualise the frequency distribution of the messages (see Figure 1(a)). At this level, interconnections between the separate displays enable the perception of an element in the context of the users' goals. Details about provenance and timeliness of information are constantly displayed, thus contributing to maintain an up-to-date awareness of the situation.

4 Context and Hierarchy Chain

Level 1 and Level 2 design recomendations can be achieved using high level overviews and drill-down capabilities. These visualisations are single-dimensional (as in they use only one feature to generate the visualisation), although interconnectable, and do not exploit the full semantic potential of the extracted features. In order to achieve Level 3 recommendations it is necessary to exploit the multiple features and provide displays that encourage establishment of correlations between data and features.

For example, for each SM message we extract a series of tags and we map them to Wikipedia entities using the DBpedia ontology. This creates a new dimension: depth, that can be combined with co-occurence to highlight the con-

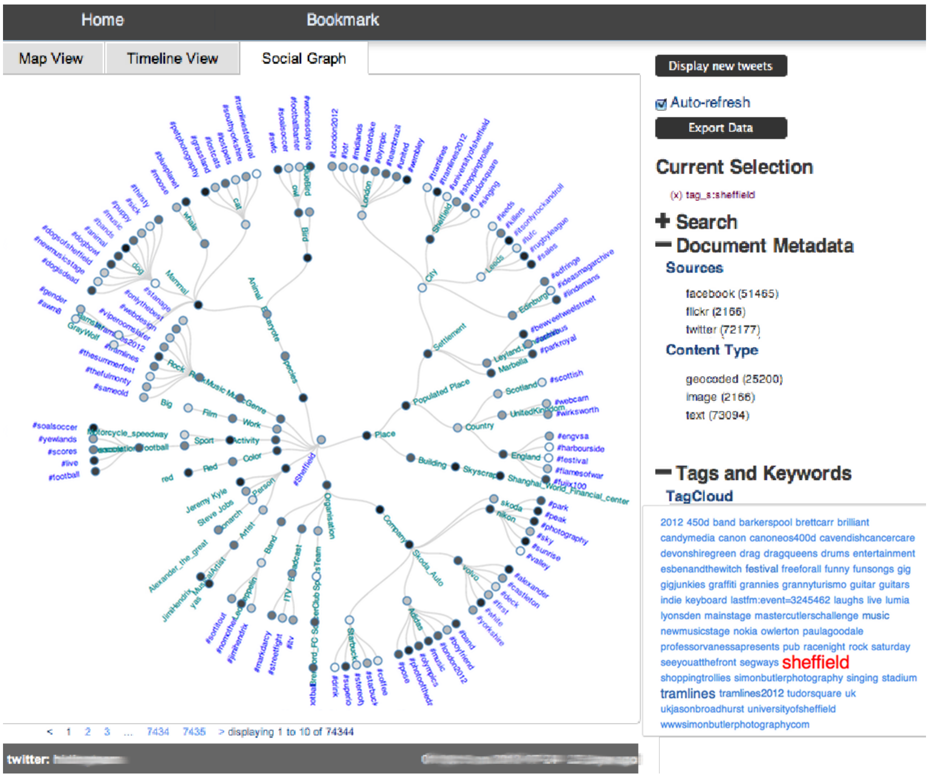

Fig. 2: An integrated framework that facilitates exploration of SM data using multiple dimensions - spatial, temporal and hierarchical. Users can select (or search for) specific tags to retrieve relevant posts, which are then visualized according to the preferred dimension of exploration - the map view, timeline view and social graph tabs on the top provides users with means to define their dimension of interest. The data shown here was an aggregation of Twitter, Facebook and Flickr posts during the 2012 Tramlines festival at Sheffield (http://www.tramlines.org.uk/)

text of the exploration as well as the ontological structure of the information. We created a new visualisation, called Context and Hierarchy Chain, that uses data features such as semantic tags and information co-occurence to provide flexible pathways to traverse and explore the information space. We define the Context and Hierarchy Chain as a hierarchical visualization technique that presents the specificity of data (to provide a multi-granular visualisation) without losing the context of the data (to provide related information)

The Context and Hierarchy chain takes inspiration from the keyword chain, a visualisation widget that is used to encourage serendipitous discovery of information in a digital library [22]. The keyword chain visualisation is designed to achieve a set of visualisation goals that are very similar to the ones of our research, in particular:

- Provide multiple visual access points
- Highlight adjacency
- Provide flexible visual pathways

The keyword chain is designed to show "relations between books based on their keywords" [22]. For each book, a series of keywords are extracted and displayed in radial curves and each of the keywords is attached to a book that shares the same keywords.

Multiple features are used to create the Context and Hierarchy chain: co-occurring tags, relatedness of tags and hierarchy of tags. These features are represented using different integrated visual metaphors - **arcs** (represent a relation between entities); **nodes** (represent an entity, i.e. a tag or concept in the hierarchy); **color-coded labels** (to differentiate between tags (blue) and concepts (green)). The chain visualisation has been integrated with the LoD methodology to allow over viewing a message or a concept from multiple perspectives, zooming in on an area of interest and using it to filter the resulting information.

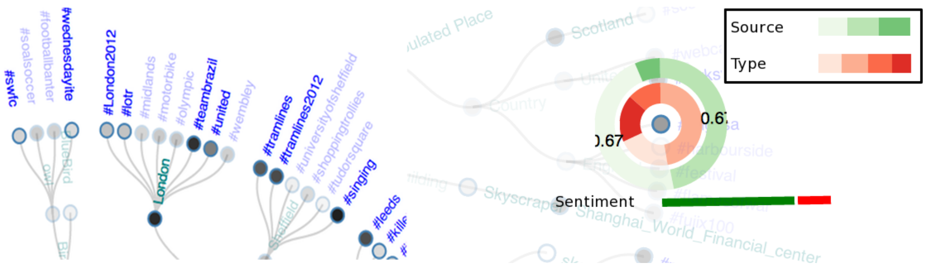

Fig. 3: **Left-** Hovering over a tag highlights related tags, greying out the rest. Here, the user hovers over the tag #tramlines2012, and the system provides the related tags - e.g. #london2012, #singing, #leeds etc.
Right- Clicking on individual elements in the context and hierarchy chain provides a summary of the information related to the element -e.g. the different types of information (text, audio, video etc), source of the information (Twitter, Flickr etc) and the overall sentiment of the posts related to the tag

The Context and Hierarchy chain plots the semantic tags associated to each concept and their parent/children accordingly to the DBpedia ontology. Concepts are attached to the central message using a hierarchical radial tree visualisation. Semantic tags (i.e. tags resolved to a concept) are extracted for each message and are attached to each concept in the hierarchy, to convey a sense of the context in which the information was transmitted. Related tags across categories are visually discoverable: when hovering on a tag, all the related ones are highlighted whilst the non-related ones are greyed out (see Fig 3, Left). In this way tags are flexibly clustered on the basis of their relatedness and different exploration paths can be followed without loosing the context of interaction.

The Context and Hierarchy chain displays the depth on the information and investigates context at different levels of specificity. For example when investigating a message related to an event in a suburb of London, it could be helpful to see related events in other London suburbs. Using the hierarchy chain, users can choose the level of specificity they desire whilst still maintaining the context of their exploration through the interlinked widgets.

Another important feature is the ability of quickly zooming in on the focus of interest and investigating it further: multiple data features have been integrated using overlays. Clicking on a tag greys out the remaining chain (but still preserves its accessibility) and a series of complementary statistical visualisations are displayed. We use a sentiment bar to present statistics about the predominant sentiment surrounding a tag, as this is a familiar metaphor for users [11]. Statistics about the source (i.e. Twitter or Flickr or Facebook) and the type of data (i.e. audio, video, textual) are displayed using a multiple donut chart, as shown in Fig. 3, Right). The donut charts are built out of sensitive sections- hovering over a section shows information related to that section (value, number of posts); clicking a section triggers a query to retrieve the relevant posts. Pressing the escape button ('esc') on the keyboard hides the overlay and enables the user to revert to their previous exploratory session.

4.1 Implementation

The context and hierarchy chain widget has been implemented as an interactive HTML widget that can be used either as a stand-alone web based interface, or as a component of other visualisation frameworks such as TUI (Tracking User Intelligence)[13]. The graph based visualization is implemented using the javascript library D3.js[7]. Hovering over items generate unique queries for looking up related tags. Clicking on items trigger queries that request for the underlying statistics of the concepts and tags. Javascript controllers then create SVG[8] elements that are overlaid on the graphical visualization, in order to present the statistics received.

5 Related Work

Information visualisation and visual analytics techniques have been often applied to semantic data, linked data and social media to enhance the user understanding of large, potentially unknown information spaces.

[7] Data-Driven Documents, http://d3js.org/

[8] Scalable Vector Graphics, http://www.w3.org/Graphics/SVG/

5.1 Semantic or Linked Data Visualisation

Based on the visual enhancement of the information being presented to the users, we group semantic systems together into several categories:

- Minimal visual enhancement (RDF stores, linked data browsers): a semantic system that mostly presents information in textual format, either as list of instances or table of triples. Examples are mSpace [21], CS-AKtive [20], PiggyBank [12]
- Moderate visual enhancement (mashup services): Systems that present data instances as interactive visual elements embedded on standard representations. Such systems typically include mashup tools that provide information to users, enriched with ad- ditional information, retrieved from linked data sources. This category also includes interfaces employing standard charting mechanisms such as pie charts, bar charts, geo- graphical map, timeline and so on to represent data. Examples are Tabulator [5] and DBpedia Mobile [4].
- High visual enhancement (visualization systems): Systems that employ a greater degree of visual encoding of data instances. Typically, systems present users with data elements that are completely converted to a representative visual element. Such visual elements are mostly highly interactive and dynamic, providing users with a greater interactive experience. Examples are IsaViz [18], RDFGravity[9], Welkin[10], Semantic Wonder Cloud [17].

An extensive survey of semantic web, and specifically LD visualisations is presented in Dadzie and Rowe [7].

5.2 Social Media Visualisation

Visual Analytics techniques have been proposed to represent and filter social media at different levels of specificity [15] [1] and to convey information evolution in the crisis management domain [19].

When visualising large scale social media data, visual analytics is mainly used to provide high level overviews. Lee [14] explores information regarding social media campaigns, Sakaki [19] uses twitter to understand the progression of earthquakes and Wongsuphasawat [23] explores trends in emergency medicine. Whilst these systems manage to efficiently display the chosen information, they are limited in the amount of data displayed. Systems with a broader focus try to capture the properties of generic data, allowing users to filter the data to items of interest. TwitInfo [15] for example, uses multiple views to present a large data set and Eddi [6] allows users to explore real time data streams relating to a given keyword. Our approach aims at combining filtering techniques as suggested by Marcus [15] and Bernstein [6] with interactive displays that enable users to determine the desired level of specificity whilst maintaining the context. We achieve this through the exploitation of semantic entities extracted from messages.

[9] http://semweb.salzburgresearch.at/apps/rdf-gravity/
[10] http://simile.mit.edu/welkin/

Whilst most social media visualisation approaches rely on geographical and temporal features TwitterReporter [16], some systems are starting to exploit the semantic of the data to enhance the visualisations. Examples of such systems are ThemeCrowds [3], TwitInfo [15], mediaWatch [11]. Marcus [15] uses features such as sentiment and link popularity to geographically plot the data. Whilst our approach provides geoplotting of social media data, we differ as we mainly focus on context and hierarchy to provide quick insight on possible interrelations between data.

mediaWatch uses features such as sentiment to create news flow diagrams that analyses the evolution of keywords and sentiments over time with an innovative display. Adams et al. [2] also focuses on interactive colour-coded timeline displays. Our approach does not focus on evolution over time yet, despite being an interesting dimension that could be exploited in Context and Hierarchy chain. ThemeCrowds [3] cluster groups of users and their evolution over time for a particular topic. Similarly to our approach, they use hierarchy to explore the data. A multi-levelhirerarchical tag cloud is built using a treemp visualisation. Our approach differentiates from Archambault [3] as we aim to provide multiple access points to the data and contextualise related data, by using a hierarchical radial tree.

6 Conclusions and Future Work

In this paper we describe a new approach to visualise contextual information from SM to support rapid exploration and analysis for non-technical users. Our research targets the Emergency Response Community, providing means for ERs to browse and use SM as an added source of intelligence through visual techniques. Information in social media is often decontextualised, due to being incomplete, informal and concerning the short-term. Whilst a user could manually correlate messages, this becomes complicated with the increasing amount of messages and the speed of arrival. To foster understanding, the pieces of information need to be automatically contextualised and correlated. Moreover, given the large scale of the information space, users must be able to focus on their interest and see the results aggregated around their focus. As our aim is to increase SA during an event, we used a state of the art definition of Levels of SA [8] translating them into design recommendations for a visual interface. The design recommendation have been implemented in a visual interface, with a particular attention for finding means to establish correlation and flexible visual pathways between data and features. We designed and developed a new visualisation approach, called Context and Hierarchy chain, that builds on semantic features extracted from the data. The Context and Hierarchy chain is used to correlate co-occurring tags or their hierarchy to provide a visual environment where data analysis can take place.

Future work will be twofold: one path will lead to the definition of Information Extraction techniques to extract more contextualised information from SM data and thier relations. Additionally, more effort will be invested in establish the semantic relatedness between both data features (topics, tags, users) and messages. Another direction of research will be the extension of the visual widgets for navigating the information space. To this extent, we will initially evaluate the existing widgets, and the Context and Hierarchy chain in particular, to understand their suitability for the Emergency Response task, by performing a user evaluation with expert personnel. The results of the evaluation will inform a re-design phase, where issues will be addressed. We will also investigate new solutions to address the design recommendations, by making use of all the additional features and relations discovered during the Data Processing step.

Acknowledgements This research was conducted as a part of the RAnDMS (Real time Analysis of Digital Media Streams) project, RC Grant Reference: EP/J020583/1, funded by the EPSRC.

References

1. Aniket Mulye Abhishek Pandya and Soon Tee Teoh. Enhancing online timeline visualizations with events and images. In *Electronic Imaging, Visualization and Data Analysis (VDA)*, 2011.
2. Brett Adams, Dinh Phung, and Svetha Venkatesh. Eventscapes: visualizing events over time with emotive facets. In *Proceedings of the 19th ACM international conference on Multimedia*, MM '11, pages 1477–1480, New York, NY, USA, 2011. ACM.
3. Daniel Archambault, Derek Greene, Pádraig Cunningham, and Neil Hurley. Themecrowds: multiresolution summaries of twitter usage. In *Proceedings of the 3rd international workshop on Search and mining user-generated contents*, SMUC '11, pages 77–84, New York, NY, USA, 2011. ACM.
4. Christian Becker, Christian Becker, and Christian" Bizer. Dbpedia mobile: A location-enabled linked data browser.
5. Tim Berners-lee, Tim Berners-lee, Yuhsin Chen, Lydia Chilton, Dan Connolly, Ruth Dhanaraj, James Hollenbach, Adam Lerer, and David" Sheets. Tabulator: Exploring and analyzing linked data on the semantic web. *In Proceedings of the 3rd International Semantic Web User Interaction Workshop*, 2006.
6. Michael S. Bernstein, Bongwon Suh, Lichan Hong, Jilin Chen, Sanjay Kairam, and Ed H. Chi. Eddi: interactive topic-based browsing of social status streams. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, pages 303–312, New York, NY, USA, 2010. ACM.
7. Aba-Sah Dadzie and Matthew Rowe. Approaches to visualising linked data: A survey. *Semantic Web*, 2(2):89–124, 01 2011.
8. MR Endsley. Theoretical underpinnings of situation awareness: a critical review. 2000.
9. GW Furnas. Generalized fisheye views. *SIGCHI Bull.*, 17(4):16–23, 1986.
10. Michiel Hildebrand, Jacco van Ossenbruggen, and Lynda Hardman. /facet: A Browser for Heterogeneous Semantic Web Repositories. In *ISWC*, 2006.
11. Alexander Hubmann-Haidvogel, Adrian M. P. Brasoveanu, Arno Scharl, Marta Sabou, and Stefan Gindl. Visualizing contextual and dynamic features of micropost streams. In *2nd Workshop on Making Sense of Microposts (MSM-2012*, 2012.

12. David Huynh, Stefano Mazzocchi, and David Karger. Piggy bank: Experience the semantic web inside your web browser, 2005.
13. Vita Lanfranchi, Rodrigo Fontenele Carvalho, Anna Lisa Gentile, Suvodeep Mazumdar, and Fabio Ciravegna. A semantic knowledge management framework for informal communication exchanges. In *Demo Track, 10th International Semantic Web Conference (ISWC 2011)*.
14. Kyumin Lee, James Caverlee, Zhiyuan Cheng, and Daniel Z. Sui. Content-driven detection of campaigns in social media. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, CIKM '11, pages 551–556, New York, NY, USA, 2011. ACM.
15. Adam Marcus, Michael S. Bernstein, Osama Badar, David R. Karger, Samuel Madden, and Robert C. Miller. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 227–236, New York, NY, USA, 2011. ACM.
16. Brett Meyer, Kevin Bryan, Yamara Santos, and Beomjin Kim. Twitterreporter: Breaking news detection and visualization through the geo-tagged twitter network. In Wei Li, editor, *CATA*, pages 84–89. ISCA, 2011.
17. Roberto Mirizzi, Azzurra Ragone, Tommaso Di Noia, and Eugenio Di Sciascio. Semantic wonder cloud: exploratory search in dbpedia. In *Proceedings of the 10th international conference on Current trends in web engineering*, ICWE'10, pages 138–149, Berlin, Heidelberg, 2010. Springer-Verlag.
18. Emmanuel Pietriga and Emmanuel" Pietriga. Doi: 10.1145/1148493.1148532 semantic web data visualization with graph style sheets. 2007.
19. Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.
20. m. c. schraefel, Nigel R. Shadbolt, Nicholas Gibbins, Stephen Harris, and Hugh Glaser. Cs aktive space: representing computer science in the semantic web. In *Proceedings of the 13th international conference on World Wide Web*, WWW '04, pages 384–392, New York, NY, USA, 2004. ACM.
21. m.c. schraefel, Max Wilson, Alistair Russell, and Daniel A. Smith. mspace: improving information access to multimedia domains with multimodal exploratory search. *Commun. ACM*, 49(4):47–49, April 2006.
22. Alice Thudt, Uta Hinrichs, and Sheelagh Carpendale. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, CHI '12, pages 1461–1470, New York, NY, USA, 2012. ACM.
23. Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 1747–1756, New York, NY, USA, 2011. ACM.