

This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

The role of visualisations in social media monitoring systems

PLEASE CITE THE PUBLISHED VERSION

http://academic-bookshop.com/ourshop/prod_3918422-ECSM-2015-2nd-European-Conference-on-Social-Media-Porto-Portugal-ISBN-9781910810316-ISSN-20557213.html

PUBLISHER

© The Authors. Published by Academic Conferences and Publishing International Limited

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This work is made available according to the conditions of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) licence. Full details of this licence are available at: <https://creativecommons.org/licenses/by-nc-nd/4.0/>

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Sykora, Martin D., Thomas Jackson, Alexander Von Lunen, Suzanne Elayan, and Ann O'Brien. 2019. "The Role of Visualisations in Social Media Monitoring Systems". figshare. <https://hdl.handle.net/2134/18757>.

The Role of Visualisations in Social Media Monitoring Systems

M. D. Sykora*, T. W. Jackson*, A. v. Lünen[†], S. Elayan* and A. O'Brien*

*Centre for Information Management, Loughborough University, United Kingdom

[†]School of Music, Humanities and Media, University of Huddersfield, United Kingdom

M.D.Sykora@lboro.ac.uk

T.W.Jackson@lboro.ac.uk

A.F.Vonlunen@hud.ac.uk

S.Elayan2@lboro.ac.uk

A.O-Brien@lboro.ac.uk

Social-Media streams are constantly supplying vast volumes of real-time User Generated Content through platforms such as Twitter, Facebook, and Instagram, which makes it a challenge to monitor and understand. Understanding social conversations has now become a major interest for businesses, PR and advertising agencies, as well as law enforcement and government bodies. Monitoring of social-media allows us to observe large numbers of spontaneous, real-time interactions and varied expression of opinion, often fleeting and private. However, human, expert monitoring is generally unfeasible due to the high volumes of data. This has been a major reason for recent research and development work looking at automated social-media monitoring systems. Such systems often keep the human "out of the loop" as an NLP (Natural Language Processing) pipeline and other data-mining algorithms deal with analysing and extracting features and meaning from the data. This is plagued by a variety of problems, mostly due to the heterogenic, inconsistent and context-poor nature of social-media data, where as a result the accuracy and efficacy of such systems suffers. Nevertheless, automated social-media monitoring systems provide for a scalable, streamlined and often efficient way of dealing with big-data streams. The integration of processing outputs from automated systems and feedback to human experts is a challenge and deserves to be addressed in research literature.

This paper will establish the role of the human in the social-media monitoring loop, based on prior systems work in this area. The focus of our investigation will be on use of visualisations for effective feedback to human experts. A specific, custom built system's case-study in a social-media monitoring scenario will be considered and suggestions on how to bring back the human "into the loop" will be provided. Also some related ethical questions will be briefly considered. It is hoped that this work will inform and provide valuable insight to help improve development of automated social-media monitoring systems.

Keywords: Social-Media Monitoring, Visualisations, User Interface Design, Decision Support Systems, Twitter

1. Introduction

Over the past years numerous social-media applications have emerged as the dominant online information sharing platforms. The content produced every day through the use of these applications is vast and carries the typical characteristics associated with big-data, such as high volume, velocity and variety in data (Laney 2001; Kitchin 2014). These characteristics complicate the automated monitoring and analysis of such datasets, especially when the requirement is to achieve (near-)real-time and scalable performance. Human driven expert processing and analysis is also unfeasible due to the volume of data; hence automated social media monitoring and analytics systems are required to deal with the social-media data overload (Fan and Gordon 2014). In this paper we present a study of social-media analytics visualisations and discuss the need and requirement for effective user interfaces and visualisations, which ought to bring the human back "into the loop" within such systems. The research presented in this paper is driven by the following three main motivations.

First of all, several researchers (e.g. Fan and Gordon 2014, Goonetilleke et al. 2014) have now highlighted the crucial importance that visualisations play within social-media analytics systems. Fan

and Gordon (2014) describe a three stage analytics system pipeline for data processing, where the final stage is concerned with the presentation of analysed data, which summarises and visualises data in order to help facilitate insight and exploratory analysis for the human expert. This visual analytics stage is described by Fan and Gordon (ibid. p 78) as *“a collection of techniques that use graphical interfaces to present summarized, heterogeneous information that helps users visually inspect and understand the results of underlying computational processes”*, which should ideally allow users to *“interactively interrogate the underlying data”*. The other two stages in their pipeline are concerned with data retrieval / collection, and automated modelling and pattern detection through natural language processing (NLP) or other data-mining techniques. Similarly, Goonetilleke et al. (2014) discuss Twitter specific social-media analytics systems at length, highlighting the need for integrated solutions which can handle the entire workflow of the data analysis life cycle, from collecting tweets to presenting the results to users. Although Goonetilleke et al. (ibid.) focus their attention on scalability of such systems to big-data, the feedback to users is indeed highlighted as an integral element in their presented system’s model. Nevertheless, Stavrakantonakis et al. (2012), in their proposed approach to evaluation of social-media monitoring tools, have also highlighted the importance of user interfaces as being one of the three main evaluation criteria for social-media systems. Many including Stavrakantonakis et al. (ibid), Goonetilleke et al. (2014) fall short in focusing on user interface visualisations’ evaluation in more detail.

Secondly, the reliability and accuracy of automated analytics is severely lacking and still leaves much to be desired. This is due to various issues with existing pattern detection and modelling techniques, further complicated by the sparse, context-poor and informal nature of social-media datasets (Ritter et al. 2011, Goonetilleke et al. 2014). Subsequently, automated analytics should arguably only be employed as decision support systems and to aid a human expert to better engage with and monitor social-media. Presenting human experts with appropriate visualisations is hence key to successful analytics systems. Furthermore, new roles in organisations and companies, such as the social-media officer, have emerged recently. Hughes and Palen (2012) for instance analysed how the traditional public information officer role has evolved to adapt to the presence of social-media, and Pacauskas et al. (2014) have even suggested that companies now establish social media departments, to provide the internal expertise and capability to deal with social-media monitoring and engagement. Given this increase of social-media professionals, there is an argument to be made for inclusion of more complex and powerful visualisations that go beyond the simplistic insights provided by line, bar, or pie charts.

Thirdly, up until now most social-media monitoring visualisations were relatively simplistic (Mathioudakis and Koudas 2010, Rogstadius et al. 2011, or MacEachren et al. 2011). For instance Mathioudakis and Koudas (2010) presented their Twittermonitor system, which automatically extracted topics from tweets to show emerging trends using quite involved data-processing techniques to achieve this. However, the associated user interface consisted of a simple view of example tweets and statistics with only basic line charts to illustrate the discovered topical trends. Such simplistic visualisations are still dominant in commercial solutions (e.g. Radian 6, Stavrakantonakis et al. 2012). To this end, Cheong and Ray (2011) discussed the use of several more involved visualisations for social-media data and argued that some of the advanced visualisation would be good candidates for future adoption in analytics solutions. Only recently several academic projects have started to emerge, employing more involved visualisations in this domain (Cheng et al. 2013, Fischer et al. 2014, or Resnick et al. 2014).

In line with this trend we introduce a set of existing yet more advanced visualisations than commonly available in the social-media analytics context which provide an evaluation of these visualisations. This is done as a case-study evaluation with human participants, using visualisations based around automated sentiment analysis of social-media messages, specifically a Twitter dataset. As far as the authors are aware this study is novel in the range of advanced visualisations for exploring social-media sentiment that are presented, and the evaluation of effectiveness of these visualisations in the

context of social-media analytics in general. The authors also followed state-of-the-art suggestions for the design and implementation of the presented visualisations.

The remainder of the paper is organised as follows. Section 2 introduces some background and prior work in the social-media analytics visualisation field, and gives brief method details, including evaluation of user interfaces and sentiment analytics which are part of our system. The custom social-media analytics visualisation system is introduced and subsequently evaluated as a case-study and the outcomes analysed in section 3. Section 4 provides a brief discussion and also highlights ethical concerns. The paper is concluded in section 5.

2. Background

Analytics systems that integrate various building blocks and elements crucial in facilitating streamlined social media monitoring with a provision for visual exploration of the analysed and underlying data have been presented in prior literature, some of which, with a focus on visualisations, are discussed in this section. The literature presented here applies mostly to Twitter; however, the work discussed can be often extended to other microblogging and related social networking platforms (e.g. Facebook, Google+...).

Osborne et al. (2014) have developed an integrated and scalable big-data analytics system for the UK government with event detection, summarisation, sentiment detection and semantic enrichment (i.e. named entity recognition, topic classification and extraction and geo-location inference) capabilities. The user interface for this system consisted of a faceted web-based interface (or faceted browsing), which allowed users to explore a collection of tweets by applying multiple filters, along multiple dimensions or facets. However, the individual visualisations employed were relatively basic – tag-clouds and line charts – except for a zoomable and interactive geo-map visualisation. Faceted browsing is relatively common among social-media analytics user interfaces (Sykora et al. 2013b) as it allows an effective way of drilling down through the available information; however, more advanced visualisations are less frequently employed, despite arguably facilitating superior insights over the explored information datasets (Fischer et al. 2014).

A couple of advanced interactive visualisations are introduced in Resnick et al. (2014) *RumorLens* analytics system for social-media rumour exploration. They employ an interactive Sankey diagram¹, as well as a network graph chart² to allow the exploration of Twitter-based rumours. Although the authors illustrate an example use of their visualisations, an actual user evaluation was not conducted.

Fischer et al. (2014) presented a visualisation interface for an analytics system which used three advanced visualisations. Using a real data case-study, similar to the evaluation study in this paper (section 3), the authors evaluated the usefulness of these visualisations, specifically the treemap view³, network graph chart and a chord diagram view⁴. Interestingly Fischer et al. (ibid.) found that users preferred the chord diagram view the least, followed by the treemap view, with the most popular being the network graph diagram, which was the most interactive of the three views.

An interactive visualisation using the “information stream river” metaphor for exploring opinion and topics flow across Twitter was introduced in Wu et al. (2014). Their visualisation combined a Sankey diagram and a density map to interactively allow users to explore overall opinion patterns at different topic levels and drill down into detail for specific patterns. The authors subsequently evaluated their system using three case-studies and semi-structured interviews with three domain experts. Again the

¹ A type of flow diagram, in which the width of the arrows is shown proportionally to the flow quantity, and can illustrate these using multiple layers. Often selecting flow arrows can bring up or highlight further details.

² A representation of a graph with nodes and edges between relevant nodes, often interactively, highlighted in some manner.

³ Allows for displaying hierarchical data by using nested rectangles, and navigating between the different hierarchy levels.

⁴ A graphical method for displaying the inter-relationships between data in a matrix, arranged radially around a circle, with the relationships typically drawn as arcs between the points.

interactivity and engagement of the visualisation was found to be one of the strongest points and the users under evaluation enjoyed its use.

Finally, it should be highlighted that it is now a standard expectation for visual analytics systems to include a landing page with overview visualisations, also known as a dashboard. MacEachren et al. (2011) or Fischer et al. (2014) are only two example systems, where interactive dashboard views are employed. In a review of the ten popular, commercial social-media analytics systems by Stavrakantonakis et al. (2012), the authors also reported dashboards being an integral element of all ten systems they've assessed.

2.1 Methodology for Visualisation Interface Evaluation

Much has been written on different visualisation analytics interface evaluation methods, and a good introduction to the existing methods commonly employed is available in Plaisant (2004). In general, it is challenging to evaluate complex visual analytics as “insight”, the major aim of knowledge discovery using visualisation analytics, is notoriously ill defined (Wijk 2013). It has been argued that case studies with experts or other capable participants, in realistic settings are often one of the few effective and reasonable ways to evaluate visual analytics interfaces (Wijk 2013, Plaisant 2004).

Similar to Wu et al. (2014) and Fischer et al. (2014), in this paper (section 3.2) we employed a case-study approach, using semi-structured interviews and several Likert-scale questions with loosely-defined tasks to gather and summarise feedback from five users, who are not co-authors of this paper. All tasks were loosely-defined as we wanted users to feel free to diverge from any prescribed tasks and to engage in more exploratory activities and potential insight generation. Participants were also encouraged to speak out, comment and report back on what they were able to see and understand while exploring the interface.

2.2 The Sentiment Analytics System

A recent, sentiment analytics social-media technique, called EMOTIVE, developed by Sykora et al. (2013a), is used in this work in order to drive our custom user-interface visualisation analytics. EMOTIVE is a system based on NLP and Ontology (semantic model) techniques, which automatically detects expressions of eight fine-grained emotions in sparse texts (e.g. Tweets). The system discovers the following range of well recognised human emotions; anger, confusion, disgust, fear, happiness, sadness, shame and surprise, but at the same time differentiates emotions by strength (also known as activation level, e.g. fear – ‘uneasy’, ‘fearful’, ‘petrified’), and among other things handles negations, intensifiers, conjunctions and interjections.

The reason this technique was chosen as the main analytics to drive our visualisations was that standard approaches to sentiment analytics tend to only perform a polar sentiment classification (i.e. negative or positive sentiment), which generally results in an output with only two dimensions. However, in (Sykora et al. 2013a) the sentiment analysis has as many as eight dimensions, hence an effective way of visualising more complex outputs is presented.

3. Case Study: Visualisations for Social Media Sentiment Analytics

In this section our visualisation analytics interface is first introduced and then evaluated. The evaluation case-study is centred around a Twitter dataset, which was collected in real-time on a major sports event (*i.e. the 2015 RBS Six Nations Rugby Championship*). In order to guarantee the same data for analysis to all participants of this study, a two week snapshot of the live dataset was used. This included the period of pre-tournament build-up, up until the 3rd of February 2015, three days before the first opening-match, and consisted of 49,713 tweets submitted by 26,258 unique Twitter

users⁵. The visualisations were designed to allow a useful, usable and engaging exploration of all tweets and emotions in the dataset (*i.e. 4,282 tweets contained explicit emotions, around 9%*).

3.1 Visualisation Analytics Interface

Our web-based user interface was implemented using PHP and HTML 5, CSS, Bootstrap, JQuery and most importantly the SVG based D3.js JavaScript visualisation library. The technology choice goes hand-in-hand with the argument put forward by Booth et al. (2014) for D3.js (Data Driven Documents) based web visualisations. Booth et al. (ibid.) make a compelling case for its use in visual analytics and provide detailed arguments for why D3.js is better than many related technologies.




Figure 1: Dashboard (overview at top, navigation and tracked hashtags on left-side, and an interactive segmented doughnut chart – a.k.a. sunburst chart)

The dashboard in figure 1 is the home-page for our analytics interface. What stands out here is the sunburst chart, which is meant to provide a very fast and relatively complete interactive overview of the dataset and help explore detected emotions. It has three levels one can hover over and drill down by topic (*e.g. Welsh, Irish or English teams*), emotion type (*e.g. happiness, anger...*) and actual expression (*e.g. looking forward to, furious...*). Clicking on the areas of the two latter layers will drill down into the individual tweets, showing these in a separate window, which is meant to further engage the users and help explore specific patterns. A separate visualisation analytics page provides the wider choice of visualisations, which can be selected by topic, and the data further constrained by date-ranges. Given that many patterns on social-media are temporal this is a useful feature to explore such patterns with. A standard stacked area chart (figure 2) is also appropriate for exploring changes over time of the eight basic emotions, and although not shown in the figure we have implemented absolute values and stream-based views as optional stacked area views.

Figure 2 highlights a Treemap and a Concentric Circles visualisation, and although these show more or less same data (*i.e. the eight basic emotions with a breakdown of actual emotional expressions*), they are visually quite different representations. The temporal Heatmap provides an overview of all tweets or tweets per each of the eight emotions, by showing an aggregated weekday x hourly time matrix.

⁵ Further details (e.g. all the 19 tracked Twitter terms) and overview statistics for this dataset are available online at the following URL – <http://emotive.lboro.ac.uk/resources/ECSM2015/>.

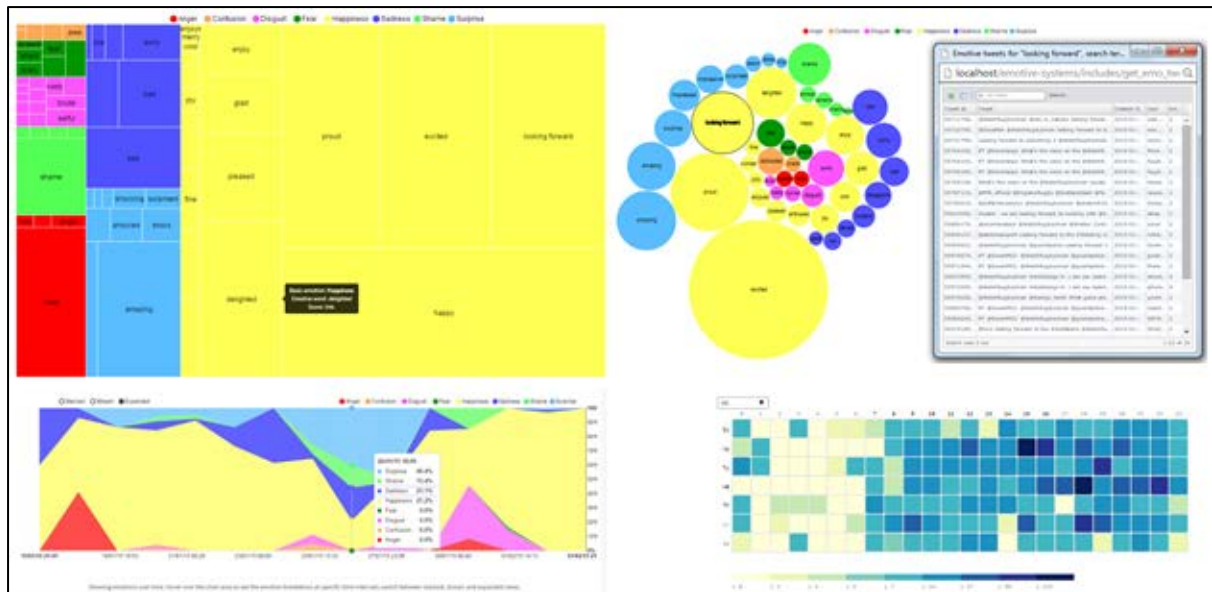


Figure 2: Clockwise from top-left; Treemap, Concentric Circles, Temporal Heatmap and Stacked-area visualisations

Figure 3 illustrates the interactive animated Bubble chart, which is useful for monitoring how emotions “bubble-up” over time. The animation starts automatically upon loading the page or selecting a topic or date-range. Hovering over the date in the bottom-right chart area will stop the animation and allow manual transition between dates. The radius of an individual circle represents the number of tweets containing a specific emotion at that point in time. The x-axis shows the cumulative number of tweets per emotion (i.e. how many tweets there are containing the respective emotion up to this point in time). The y-axis shows the intensity of the emotive term, i.e. the higher the circle is up on the y-axis, the higher the emotive score is. Therefore, this view is particularly useful for monitoring the rate of change in emotional tweet content activity.

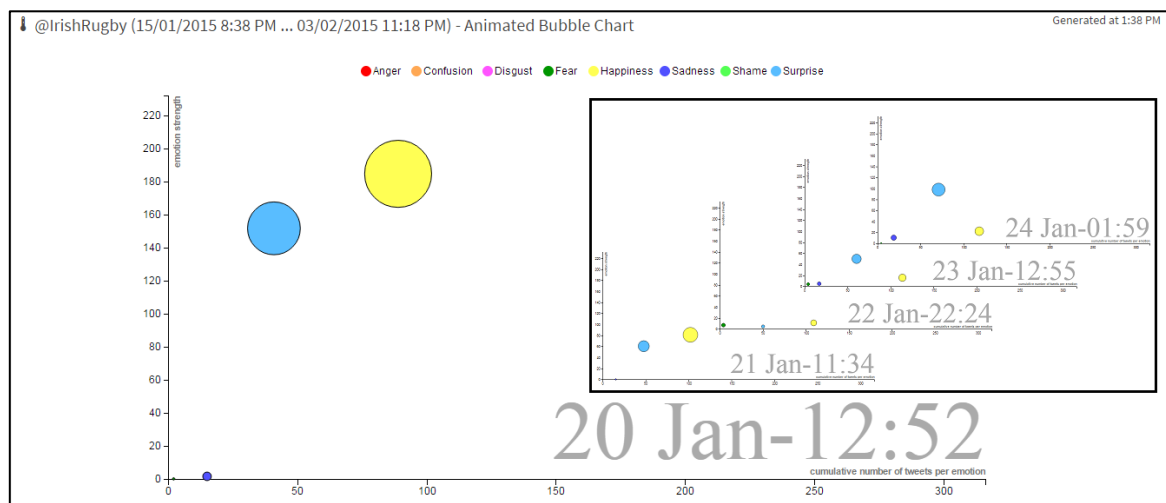


Figure 3: An animated Bubble chart showing the eight emotion dimensions evolving over time – the inset figure highlights example changes over time

The navigation-menu also allows bringing up a Gallery of top emotional Twitter users and also the most recent images shared on the various topics. Both views are presented as interactive walls, with a dynamic, per emotion type, tile display and a brick-wall gallery view with a possibility of image slideshows, respectively.

3.2 Evaluation Study

In total five participants who haven't used the visualisations before (*i.e.* 4 postgraduate students in information science and one research associate, none of who was a co-author of this paper) participated in the case-study evaluation. On average each session took around 50 minutes. Our main focus was to collect a qualitative assessment of the visual analytics system and the related user interface, as well as attempting to quantitatively evaluate the user experience of the interactive visualisations in terms of familiarity, usefulness and engagement⁶.

Users were asked to first explore the dashboard, navigate to the visualisations page and explore the dataset / emotions analytics using the available set of visualisations, and finally users were expected to browse through the user and images galleries. The four advanced visualisations, *Treemap*, *Concentric Circles*, *Temporal Heatmap* and *Animated Bubble chart*, were of most interest to us, as these are still all relatively non-standard in social-media visualisation analytics. Considering that all five participants were educated to at least a degree level in the information sciences field, it was interesting to observe that prior to using the visualisations their familiarity with the four advanced visualisations, on a five point Likert-scale, only scored an average value of 2.4 (standard deviation 1.28). However, despite high unfamiliarity, the average score for all five users in terms of usefulness and engagement were 4.1 (std. dev. 0.83) and 4.2 (std. dev. 0.81), respectively.

Prior to its use, users were by far the least familiar with the animated Bubble chart (average 1.4, standard deviation 0.55), but were the most familiar with a Treemap interface. Concentric circles and Heatmap charts scoring in-between (for all, see Avg. and Std. Dev. columns in table 1). Concentric circles scored the highest in terms of their perceived usefulness (avg. 4.6, std. dev. 0.55). Interactivity was commented on by several participants as a desired feature, and it's hence interesting to take a look at the scores of perceived engagement of the various visualisations. As expected, the animated Bubble chart had one of the highest engagement scores (avg. 4.6, std. dev. 0.55). Also, although the Treemap and Concentric circles both show very similar data, Concentric circles had a much higher engagement score (avg. 4.6 vs. avg. 3.6), despite users' overall lower familiarity with this particular visualisation. User 1 for instance spent most of their time exploring data through this particular visualisation and clearly enjoyed its interactivity. User 2 commented that some visualisations (especially in relation to Concentric circles), look like "an artistic picture and have a beautiful look about them" making it "pleasant to interact with".

Table 1: Treemap, Concentric Circles, Heatmap and the overall User Interface Likert ratings (5-point scale) for Familiarity, Usefulness and Engagement.

	User 1	User 2	User 3	User 4	User 5		Avg.	Std. Dev.
Treemap Familiar	3	1	3	5	4		3.2	1.48
Treemap Useful	4	5	4	4	4		4.2	0.45
Treemap Engaging	4	4	3	5	2		3.6	1.14
ConcCirc Familiar	3	2	1	2	4		2.4	1.14
ConcCirc Useful	4	4	5	5	5		4.6	0.55
ConcCirc Engaging	5	5	4	4	5		4.6	0.55
Heatmap Familiar	2	2	1	5	3		2.6	1.52
Heatmap Useful	2	3	4	5	4		3.6	1.14
Heatmap Engaging	3	4	4	5	4		4	0.71
AnimBub Familiar	2	1	1	1	2		1.4	0.55
AnimBub Useful	4	5	3	3	5		4	1.00
AnimBub Engaging	4	5	4	5	5		4.6	0.55
UI Usable	3	4	4	5	4		4	0.71
UI Useful	4	5	4	5	5		4.6	0.55
UI Engaging	4	4	4	4	4		4	0.00

⁶ Defined as simply a self-assessed amount to which a particular visualisation engages a user, and invites or encourages him or her to interact with the particular visualisation.

Although user 1 found the stacked area chart (for which Likert scores were not collected) useful for exploring temporal patterns in tweet emotions, but their favourite was the animated Bubble chart though, stating that it provided "the best overview of the temporal emotional timeline". User 5 also stated that they preferred the animated Bubble chart the most, and user 2 pointed out that it was "definitely most useful, as it flags up emotions" and that "the animation is a bit of fun and entertaining to interact with too". The lowest score for usefulness was assigned by user 1 to the temporal Heatmap, although the user commented that data spanning a longer time-period would make this view likely very useful. Users 3 and 4 on the other hand were particularly vocal about the benefits and the potential for emotional weekly, daily and time of day patterns emerging from the temporal Heatmap, and pointed out its potential in areas of business and especially marketing.

Users also commented on the images gallery, which they found useful in getting a taste for the event in terms of pictures shared. More enthusiastic comments were related to the user gallery, which as for instance user 2 commented, "can be very important in highlighting highly emotional tweeters and then being able to engage with them directly".

The dashboard view itself (figure 1) was well received, highlighted by comments such as "quite neat, I really like it a lot", or "it's nice and very pretty with an intuitively familiar look". However, comments regarding the sunburst chart were mixed. Four of the users commented that they initially perceived it negatively, as it took some fiddling around to get used to the interactivity of it, and only after several minutes of familiarising themselves with it, did they comment on its usefulness in being able to quickly drill down to data behind the different topics, emotions and expressions. Being able to do such an amount of "data-zooming", right-away from the dashboard screen, was commented upon highly positively by every participant. Nevertheless advanced and hyper-interactive visualisations, such as the sunburst, likely requires further refinement and we plan to conduct future work on improving and subsequently re-evaluating these visualisation in future evaluations.

The average Likert scores for the overall user interface (last three rows in table 1), highlight the relative success as usability, usefulness and engagement of the user interface all ranked a score of 4 or higher.

4. Discussion

Many organisations have now taken to social-media analytics, to engage with and monitor social-media, and are increasingly keen on developing the capability of collecting, storing and analysing social media data for the purpose of harvesting information and actionable knowledge for decision making and forecasting. Social-media based visualisation analytics would seem important in the monitoring task. As was discussed earlier (section 2), only recently did more advanced visualisations start appearing in social-media analytics, and their full uses and appropriate evaluations still require more research. In this study it was attempted to address this lack of research and our evaluation helped to illustrate the general usefulness and importance of interactivity in visualisations analytics in this field, and our results to this end are in line with findings reported by Wu et al. (2014). More general comments from participants in the evaluation case-study hinted that users (i) like to see and relate to underlying data (i.e. by being able to drill down directly to the underlying tweets – this is possible in most of our visualisations); and (ii) being able to make relative comparisons (i.e. to compare certain values / patterns to related patterns with slightly different parameters – this is possible only in some of the views, such as the temporal Heatmap). In terms of future improvements it would be interesting to address (ii). However, there are significant ethical concerns over ownership and data rights of social-media generated datasets. Puschmann and Burgess (2014) provide an extensive introduction to these issues and suggest that ethical guidelines and contracts between platform providers and users are still in the process of maturing and although accessible now, might become a barrier to conducting certain types of research projects in future. The current situation is best described with a quote from Puschmann and Burgess; "while the data in social media platforms

is sought after by companies, governments, and scientists, the users who produce it have the least degree of control over ‘their’ data” (ibid. p 44). Social analytics and powerful interactive visualisations place great power into the hands of the few, and this in itself raises ethical questions. Finally it must be noted that all kinds of data, especially complex and multidimensional datasets, common on social-media platforms, are suitable for the here presented visualisations.

5. Conclusion

An extensive discussion of social-media visualisation analytics was undertaken in this paper. A custom set of visualisations was developed and presented with a subsequent case-study evaluation on a group of users over a real Twitter-based dataset. Our main argument proposed in this article was that there is a need for more work looking specifically at social-media based visual analytics, and that tools such as the one presented in this study are important for effective analytics feedback to human experts. In order to improve usefulness and engagement of visualisations, these need to be evaluated in the first place. We hope that the results of this study and reported experiences will be of use to others, and especially in the field of social-media monitoring, addressing how best to bring the human back “into the loop”.

References

- Booth, P., Hall, W., Gibbins, N. and Galanis, S. (2014), Visualising Data in Web Observatories: a proposal for visual analytics development & evaluation, 23rd International conference on World Wide Web, Seoul (Korea).
- Cheng, D., Schretlen, P., Kronenfeld, N., Bozowsky, N. and Wright, W. (2013), Tile based visual analytics for twitter big data exploratory analysis, IEEE International Conference on Big Data, Santa Clara (USA).
- Cheong, M. and Ray, S. (2011), A literature review of recent microblogging developments. Technical Report, Victoria, Australia: Clayton School of Information Technology, Monash University.
- Donath, J., Dragulescu, A., Zinman, A., Viégas, F. and Xiong, R. (2010), Data portraits. In ACM SIGGRAPH 2010 Art Gallery, pp. 375-383.
- Fan, W. and Gordon, M. D. (2014), The Power of Social Media Analytics, Communications of the ACM 57 (6), pp. 74–81.
- Fischer, F., Davey, J., Fuchs, J., Thonnard, O., Kohlhammer, J. and Keim, D. A. (2014), A visual analytics field experiment to evaluate alternative visualizations for cyber security applications. EuroVis International Workshop on Visual Analytics, Swansea (UK)
- Goonetilleke, O., Sellis, T., Zhang, X. and Sathe, S. (2014), Twitter analytics: a big data management perspective, ACM SIGKDD Explorations Newsletter 16 (1), pp. 11-20.
- Hughes, A., L. and Palen, L. (2012), The evolving role of the public information officer: An examination of social media in emergency management. Journal of Homeland Security and Emergency Management 9 (1), Last Accessed 18 January 2015, doi:10.1515/1547-7355.1976
- Kitchin, R. (2014), The data revolution: Big data, open data, data infrastructures and their consequences. Sage, London.
- Laney, D. (2001), 3D Data Management: Controlling Data Volume, Velocity and Variety. Available at: <http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>, last accessed: 19 December 2014.
- MacEachren, A., Jaiswal, A., Robinson, A., Pezanowski, S., Savelyev, A., Mitra, P., Zhang, X. and Blanford, J. (2011), SensePlace2: GeoTwitter analytics support for situational awareness. In Proceedings of the Visual Analytics Science and Technology (VAST) IEEE Conference, Providence (USA).
- Mathioudakis, M. and Koudas, N. (2010), Twittermonitor: trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, Indianapolis (USA).

Pacauskas, D., Durgam, P., & Fomin, V. V. (2014). How Companies Can Modify R&D for Integrating Social Media Activities into the New Products Development, eConference eEcosystems Conference, Bled (Slovenia)

Plaisant, C. (2004), The challenge of information visualization evaluation, ACM Working conference on Advanced Visual Interfaces, Bari (Italy).

Resnick, P., Carton, S., Park, S., Shen, Y. and Zeffer, N. (2014), RumorLens: A System for Analyzing the Impact of Rumors and Corrections in Social Media, Computational Journalism Conference, New York (USA).

Ritter, A., Clark, S., Etzioni, M and Etzioni, O. (2011), Named Entity Recognition in Tweets: An Experimental Study, Conference on Empirical Methods in Natural Language Processing, Edinburgh (UK)

Stavrakantonakis, I., Gagiou, A. E., Kasper, H., Toma, I. and Thalhammer, A. (2012), An approach for evaluation of social media monitoring tools, *Common Value Management*, 52 (1), pp. 52-64.

Sykora M., Jackson T. W., O'Brien A. and Elayan S., (2013a), Emotive Ontology: Extracting Fine-Grained Emotions from Terse, Informal Messages, *International Journal on Computer Science and Information Systems* 8 (2), pp. 106-118.

Sykora, M., Jackson, T. W., O'Brien, A. and Elayan, S. (2013b), National Security and Social Media Monitoring: A Presentation of the EMOTIVE and Related Systems. *Proceedings of the European Intelligence and Security Informatics Conference (EISIC)*, Uppsala (Sweden)

Osborne, M.; Moran, S.; McCreadie, R.; Von Lunen, A.; Sykora, M.; Cano, E.; Ireson, N.; Macdonald, C.; Ounis, I.; He, Y.; Jackson, T.; Ciravegna, F. & O'Brien, A. (2014), Real-Time Detection, Tracking, and Monitoring of Automatically Discovered Events in Social Media, in 'Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations', Association for Computational Linguistics, Baltimore (USA).

Wu, Y., Liu, S., Yan, K., Liu, M. and Wu, F. (2014), OpinionFlow: Visual analysis of opinion diffusion on social media, *IEEE Transactions on Visualisation and Computer Graphics* 20 (12), pp. 1763-1772.