

PREPRINT – To be published in *Aslib Journal of Information Management*, vol.67, no.3 (2015)

Public scientific communication on Twitter: visual analytic approach

Victoria Uren (Aston University) & Aba-Sah Dadzie (University of Birmingham)

Purpose

The purpose of this paper is to assess high-dimensional visualisation, combined with pattern matching, as an approach to observing dynamic changes in the ways people tweet about science topics.

Design/methodology/approach

The high-dimensional visualisation approach was applied to three scientific topics to test its effectiveness for longitudinal analysis of message framing on Twitter over two disjoint periods in time. The paper uses coding frames to drive categorisation and visual analytics of tweets discussing the science topics.

Findings

The findings point to the potential of this mixed methods approach, as it allows sufficiently high sensitivity to recognise and support the analysis of non-trending as well as trending topics on Twitter.

Research limitations/implications

Three topics are studied and these illustrate a range of frames, but results may not be representative of all scientific topics.

Social implications

Funding bodies increasingly encourage scientists to participate in public engagement. As social media provides an avenue actively utilised for public communication, understanding the nature of the dialogue on this medium is important for the scientific community and the public at large.

Originality/value

This study differs from standard approaches to the analysis of microblog data, which tend to focus on machine driven analysis of large-scale datasets. It provides evidence that this approach enables practical and effective analysis of the content of midsize to large collections of microposts.

Keywords: High-dimensional Visualization, Visual Analytics, Microblogs, Public Engagement with Science, Science Communication, Social Media Metrics.

Introduction

The current ethos of science communication, driven by policy (POST, 2003), advocates engaging the public in dialogue (Nisbet and Scheufele, 2009). As a consequence, research funding bodies typically require public engagement and dissemination of findings beyond scholarly communities (see, .e.g, the Research Councils UK's strategy for public engagement at <http://www.rcuk.ac.uk/pe>, and the European Commission's policy on public engagement at <http://ec.europa.eu/research/swafs/index.cfm?pg=policy&lib=engagement>). Social media are the latest in a series of technologies which have shaped scientific communication and enabled scientists to disseminate their research more widely (Hogan and Sweeney, 2013). It has been noted (Meyer and Schroeder, 2009) that online access to scientific discussion gives the public access to scientific information which has not been mediated by the traditional communication professions of science journalism, documentary film-making, etc. Journalists have thus also had to adapt to compete with other communicators exploiting the new media (Fahy and Nisbet, 2011; Hermida, 2010). Events related to scholarly content on social media are being explored as potential indicators of research impact (Priem *et al.*, 2012), and research communities have embraced social media for informal communication (Ponte and Simon, 2011; Darling *et al.*, 2013; Bar-Ilan *et al.*, 2012; Weller *et al.*, 2011; Mandavilli, 2011). For science organizations, Web 2.0 has become integral to public relations, and means are being sought to quantify its impact at an organizational level (Roemer and Borchardt, 2013). Individual scientists are also aware of the public engagement agenda: Letierce *et al.* (2010) have shown that although researchers' main

motivation for tweeting is to communicate with members of their own community (89%), some also try to reach general audiences (45.9%). As a source of altmetrics about formal publications, Twitter has low coverage compared to other social media such as Mendeley (Zahedi *et al.*, 2014). However, scientists play only one part in the bigger picture – science organizations, journalists, lobbyists and the general public also have important contributions to make. We find microblogs, specifically Twitter, of interest because of the low barrier to entry compared to other social media such as blogs and discussion forums. They also have the advantage of being an open forum, including people who would not normally, or ever, read scholarly literature.

Bauer *et al.*, (2007, p. 90) have called “*to expand the range of data ‘officially and legitimately’ relevant for monitoring public understanding of science*”. The measurement of science communication in Web 2.0 media, which is sometimes known as altmetrics (Priem *et al.*, 2012), plays a role here. The dominant strand in social media metrics research, so far, is in assessing impact of scholarly papers on social media (Haustein *et al.*, 2014a, 2014b; Shuai *et al.*, 2012; Thelwall *et al.*, 2013; Eysenbach, 2011). Such work extends previous bibliometric studies on the communication of scholars. In addition to “*evaluation of scholars*” and “*recommendation of articles*” Priem and Hemminger (2010) define the remit of the scientometric study of Web 2.0 to include “*the study of science*”. It is at this point, at the intersection of the science communication studies and social media metrics, that this work positions itself.

In the study presented here, we have observed relatively informal science communication on Twitter. We used selected scientific terms to filter samples of

science related communication by any Twitter account holder. Our aim was to understand better which aspects of science provoke transient and sustained activity on the microblogging platform. The work, therefore, focused on identifying ‘message frames’ within the tweets, where frames are seen as ways of interpreting topics, identifiable by the use (or avoidance) of certain words and phrases. The study applies a mixed methods approach to framing analysis for science communication in social media, combining content analysis with high-dimensional visualisation of frames identified using pattern matching. The contribution of the work lies in the assessment of the visualisation of science related tweets as a method for analysing communication and improving the understanding of Twitter’s role in the context of altmetrics.

The aim of this study was to provide proof of concept using the new method, which combines content analysis and visual analysis as a means to observe dynamic changes to the framing of science communication on Twitter: which frames were prominent when and how this changed over time. One specific issue facing Twitter research on *scientific* topics is that science produces relatively few posts compared to current affairs, popular culture, etc. (Uren and Dadzie, 2011). This accurately reflects the lower interest in these topics by lay society as a whole, but nonetheless makes gathering usable samples difficult. On these small samples, simple trend spotting methods can be ineffective. Our first research question addressed the scale issue:

1. Does the proposed method support the analysis of dynamic changes in non-trending topics?

A requirement of a method for observing science communication is the ability to support longitudinal studies, in order to facilitate medium to long-term observation.

This is especially important for microblogging services such as Twitter, to ensure that spurious spikes that do not represent truly trending topics do not skew the analysis.

Therefore our second research question was:

2. Can changes be observed across disconnected time periods (within days and in samples taken a year apart)?

The established method for this kind of study would be content analysis. However, content analysis requires multiple rounds of analysis and cross checking between coders, which takes time. Further, only relatively small tweet counts can be handled this way, which could become problematic for trending topics or long-term longitudinal studies. We therefore combined content analysis with high-dimensional visualisation of tweeting activity, using parallel coordinates (Inselberg, 2009); the content analysis takes a subjective approach to the identification of frames in smaller samples, whereas the visual analytics deploys pattern matching to explore communication in larger datasets. This approach allowed us to compare multiple frames (coordinates or dimensions) across multiple time periods, and where sufficient data was available, at larger scale. A third question, aimed at determining whether the addition of a visual analytic approach added value to the content analysis was:

3. Does visualisation-based analysis reveal further information in addition to confirming the content analysis?

Literature Review

We have identified several studies, which address public scientific communication in microblogs. Some are topic focused, considering issues such as climate change (Hubmann-Haidvogel *et al.*, 2012), nanotechnology (Veltri, 2013), and astronomy

(Wilkinson and Thelwall, 2012). Others look at the dispersion of messages (Chew and Eysenbach, 2010; de Domenico *et al.*, 2013), while a number focused on Twitter activity around events (Adams *et al.*, 2011; Desai *et al.*, 2012). These publications confirm interest in analysing scientific microblogs and point to a diversity of analytical methods, with the results targeted at audiences with different perspectives.

The content of tweets has been studied using a range of natural language processing (NLP) methods. The 140 character restriction on tweets poses challenges here, leading to unconventional grammar and abbreviations, although language use on Twitter has been shown to be surprisingly formal (Hu *et al.*, 2013). Twitter specific NLP methods have been developed, for example, Reddy Yerva *et al.* (2011) present a classification method for ambiguous terms (e.g., the product names Apple and Blackberry). Ritter *et al.* (2011) present a named entity recognition (NER) algorithm, and Thelwall *et al.* (2011) developed SentiStrength, a sentiment analysis system, all for short social web texts. Because of the important temporal element in Twitter analysis, topical content analysis of tweets is frequently aligned with events.

Exploration of topic evolution on social media has been tackled by several methods: Hu *et al.* (2012) use a Latent Dirichlet allocation (LDA) based method, Cui *et al.* (2011) use a hierarchical Dirichlet process, Dou *et al.* (2013) similarly employ hierarchical text analytics with interactive visualization, Chua and Asur (2013) have developed a method for summarising Twitter content by extracting representative tweets.

In contrast to the computational approaches for text analysis outlined above, content analysis employs human coders to classify texts – defined as a qualitative “*research*

technique for making replicable and valid references from texts” (Krippendorff, 2004). Although much Twitter analysis has focused on trending topics with millions of tweets and on getting results quickly, (manual) content analysis has been employed by researchers seeking subtle information about content which goes beyond current capacity of (automated) NLP methods for these short, irregular texts. Examples of content analysis applied to tweets include: Chew and Eysenbach (2010), in which a sample of tweets about the H1N1 flu pandemic was categorized for content (e.g., resources, personal experience), qualifiers (e.g., humour, relief, concern) and links; and analysis of discourse on stem cell science (Adams *et al.*, 2011) for tone (sentiment), user identity, and a range of frames such as religious, political and business.

In this work we present an approach which bridges qualitative content analysis and NLP by using the former to guide deployment of the latter. We do this through the use of a visual analytics method. Visual analytics has evolved as a field where interactive visualization is used to augment human intuition and perception during analysis of complex data. First, high-level, exploratory overviews of data structure and content are examined. Detailed analysis of regions of interest (ROIs) thus discovered is then carried out using visualization methods selected based on data type and task, and the target audience.

The most basic visualization used in Twitter analyses is the timeline plot of numbers of tweets, trends, etc., (e.g., Lumezanu *et al.*, 2012; de Domenico *et al.*, 2013). Word clouds (e.g., Xu *et al.*, 2013), and variations of scatter plots (e.g., Cao *et al.*, 2012; Yuan *et al.*, 2009) are other simple options used to examine term usage. Tweets often

include explicit, or automatically derived, geo-location information: another dimension along which to categorize them. The results of such analysis are typically displayed using cartographic or other location-based visualisation, topic distribution (Cano *et al.*, 2013) or evolution (de Domenico *et al.*, 2013). Analysis of the evolution of topic and sentiment in dynamic datasets requires techniques able to handle multi-attribute data; examples include flow visualizations such as streamlines or themerivers, that harness the metaphor of a flowing river (e.g., Cui *et al.*, 2011; Dou *et al.*, 2013; Xu *et al.*, 2013).

Visualization also has a long history in the field of bibliometrics (Börner *et al.*, 2003), where it has a role in mapping knowledge domains and their development over time. The dominant visualisation maps topic clusters in 2D space as used in VOSviewer (van Eck and Waltman, 2010). Time may be represented with a view for each time window, e.g., as by Kraker *et al.* (2014) and in the cluster network visualization of the SciMAT tool (Cobo *et al.*, 2012). Other visualizations present citation networks plotted against a time axis, as in CitNetExplorer (van Eck and Waltman, 2014) or the longitudinal view in SciMAT.

Parallel coordinates provide another method for visual comparison of multiple, dynamic data attributes. In this work we deploy them to visualize multiple time windows and multiple frames in the same view, allowing the observation of changes in interest and terminology usage through time. The remainder of this paper demonstrates the application of parallel coordinates as an analytical tool for the selected scientific topics, to answer the research questions outlined above.

Methods

This study observed one trending topic (the Curiosity landing) to represent a large-scale event, and two non-trending topics (Phosphorus and Permafrost), which were identified as topics in which previous research showed low, but consistent, Twitter activity (Uren and Dadzie, 2011). Non-trending communication about science can be assumed to involve what Miller (1983) characterises as *attentive* groups: fewer in number, but persistently engaged. The samples provide contrast, allowing us to investigate whether methods for studying trending topics transfer to non-trending, but not less important, topics.

A longitudinal study was carried out with two sampling periods: 4-9 August 2012 (the period of the Curiosity landing) and 4-9 August 2013 (its anniversary). Data was collected using the Twitter streaming API for the terms: “curiosity”, “phosphorus” and “permafrost”. We did not sample using hashtags specifically, as many studies do, because of the relatively low numbers expected. The non-bounded pattern search for, e.g., “curiosity”, however, also retrieves #curiosity. Subsequently, Tufekci (2014) has written concerning social sampling biases of hashtag studies, which supports our decision not to use them here. Total tweets sampled for each period were:

- Curiosity 2012 – 1,194,470
- Curiosity 2013 – 3,310
- Phosphorus 2012 – 587
- Phosphorus 2013 – 6,269
- Permafrost 2012 – 311
- Permafrost 2013 – 618

Curiosity, also known as MSL (Mars Science Laboratory) and on Twitter as the persona @MarsCuriosity, is a robot designed to conduct geological survey work on the surface of Mars. It landed on Mars on 6th August 2012. Our first sampling period includes the landing date. Considerable publicity surrounded the landing, heightened by the history of failed landings. There was however overwhelmingly positive response as the first images came back to earth; the Curiosity dataset provides an example of a topic for which there are many tweets representing a big international event.

Phosphorus is essential for all life. As a result, a fair number of biology articles concerning phosphorus are announced on Twitter. The biggest commercial use of phosphorus is in fertilizers (Neset and Cordell, 2012), but mineral sources are limited (“peak phosphorus”). Phosphorus fertilizers can pollute water supplies, and this can become politicised. A more active political issue, in the periods we sampled, concerns the use of white phosphorus in warfare. This bears similarities to other debates about technologies perceived as risky, such as nuclear power. The combination of these threads causes a steady, albeit relatively low, flow of tweets containing the term ‘phosphorus’.

Climate change has been extensively studied in the public engagement of science literature because of its importance and its controversial nature (e.g., Hubmann-Haidvogel *et al.*, 2012; Nisbet, 2009). The term ‘permafrost’ allows us a small view onto this wider debate. Permafrost is a soil type characterised by a permanently frozen layer (Schaefer *et al.*, 2012). Permafrost can be viewed both as an indicator of

warming (during the second, 2013, sampling period it was melting at an unprecedented rate), and a climate risk (melting releases methane). Furthermore, changes to underlying soil structure disrupt ecosystems and the built environment of affected areas. Therefore it is of scientific interest to those interested in climate change, and has political, economic, and social importance for people living in the affected regions.

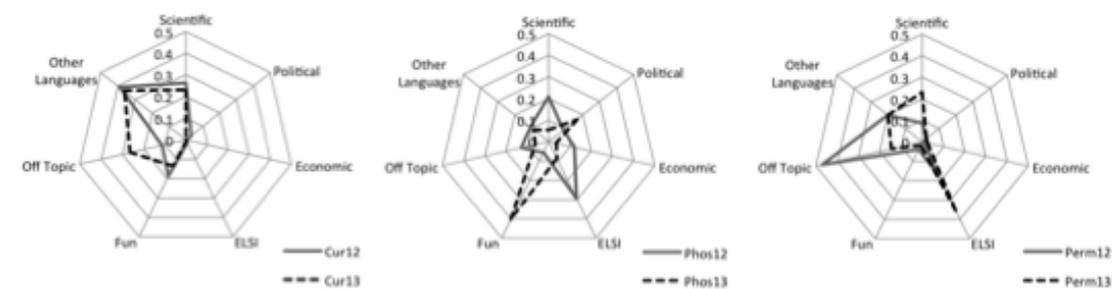
Frames

The notion of “frames”, first proposed by Goffman (1974), is influential in the interpretation of public communication about science, with frames seen as a tool for shaping public perceptions (Nisbet and Scheufele, 2009). Frames may be defined as: “*ideal-type arguments used to interpret certain topics*” (Schäfer 2009, p. 485), or “*interpretative storylines*” (Nisbet and Scheufele, 2009, p. 1770) and are marked by “*patterns in the use of certain words, phrases, images, and sources of information*” (Bruu Carver *et al.*, 2013, p. 9). Hence, frames can be recognised by the use (or avoidance) of certain words, phrases, metaphors or images, and the status of particular kinds of information sources, which together point at a particular interpretation of an issue.

For the content analysis, we sought a general-purpose framing scheme which could be applied to a range of science-related topics. The scheme proposed by Schäfer (2009) was selected because: 1) it has been applied to multiple topics (stem cell research, human genome and neutrino research), and 2) it was the simplest such scheme we

found (given the constraints of 140 characters, complex arguments are difficult to frame in tweets).

The scheme was originally used in the analysis of elite German newspapers, therefore some modifications were needed to adapt it for the less formal discourse of Twitter. To the four original codes (*Scientific, Political, Economic, ELSI – Ethical, Legal & Social Implications*) a fifth code was added to reflect the light-hearted nature of many posts on Twitter, which was labeled *Fun*. Some further codes were required to classify the tweets that could not be assigned to a frame. Firstly, scientific terms are quite international in their usage and “English” words therefore occur in tweets written in a wide range of languages (Spanish, German and Indonesian were found, among others, often also containing English terms). We determined that interpreting the subtleties of framing in 140 character tweets that had been passed through automatic translation was too error prone. Therefore, the category *Other Languages* was added. A benefit of this code is that it indicates international interest. Finally, the code *Off Topic* was defined for tweets that could not be categorised with any other code. For example, #curiosity is used in its general sense and Permafrost is the name of a games server. The results of the coding are presented in Figure 1. A coding manual providing definitions can be found at: <http://bit.ly/1s9P9DV>. The manual was refined by agreement between the two coders prior to final coding.



**Figure 1 Proportion of frame codes per dataset, left to right Curiosity,
Phosphorus and Permafrost.**

We define the unit of analysis as a single tweet. Each tweet was assigned exactly one code. For each topic, 12 batches of tweets were coded, one for each day 04-09 Aug. in 2012 and 2013. For Phosphorus and Permafrost the numbers of tweets per day were relatively low (32-116 for permafrost, 65-4,736 for phosphorus). For Curiosity, the numbers of tweets were too great to be coded in their entirety (at its maximum on the 6th of Aug. 2012 there was a total of 733,000 tweets). In 2013 Phosphorus had some days with several hundred tweets and one with over 4,000. Therefore, each batch contained the smaller of either the total number of tweets filtered on the term for the day or a random sample of 200 of the filtered tweets (randomisation was achieved using the SQL ORDER BY RAND() statement). Table 1 summarises the numbers coded, and Figure 1 summarises the frames identified for each dataset.

Dataset	Cur12	Cur13	Phos12	Phos13	Perm12	Perm13
4 Aug	200	200	94	169	70	70
5 Aug	200	200	65	200	33	108
6 Aug	200	200	113	200	55	116
7 Aug	200	200	137	200	73	114
8 Aug	200	200	86	194	49	99
9 Aug	200	192	92	200	32	111
Total	1200	1192	587	1163	312	618

Table 1 Numbers of tweets coded

Agreement between the two coders was calculated using Hooper's measure (1965):

$$H = \frac{C}{A+B-C}$$

where C is the number of codes on which both coders agree, A is the number of codes assigned by coder A, and B is the number of codes assigned by coder B. Note that Cohen's Kappa (1960) is inappropriate for this study because tweets are not independent of each other because of retweeting. Hooper's measure was chosen as an alternative because it is a well known consistency measure in document indexing, e.g., Medelyan and Witten, (2006), which has procedural similarities to *a priori* coding. Furthermore, it gives consistently lower values than the other well known alternative Rolling's measure (Leininger, 2000). The batch agreement was calculated as the mean of the agreement values per day. Frame agreements were calculated as raw values cumulated for all six days. Several rounds of coding were undertaken for the datasets. Coding was stopped when no batch in the dataset had an agreement of less than 0.7. To achieve this, coders had to reach a similar or higher level of agreement on the most populous frame codes for the dataset (some individual frames may have low agreement at this point but these would be for barely populated frames with little influence on the overall score).

Visual Analytics

Visualisation and analysis are often treated as distinct steps. However, increasingly, (interactive) visual analytics is being used to tackle more effectively the analysis of very large amounts of complex data, by augmenting advanced human perception with

high computing power (see, e.g., Cao *et al.*, 2012, Cui *et al.* 2011, Xu *et al.*, 2013). For our study, the qualitative visualisation, while human-driven, relies on the capability of the visualisation technique, parallel coordinates (Inselberg, 2009), to highlight trends within the data, and therefore reveal ROIs, key here being areas of intense activity and other areas where unexpected patterns occur, to determine where to investigate in more detail. Where ROIs were identified, further detailed analysis of tweet content was carried out, by investigating manually the tweets used in the content analysis (see Table 1), to identify additional terms or evolution in terminology usage, in addition to the larger samples visualised. The visual analytics approach was thus used to guide the identification of terms tweeters use to express their opinions about and interest in different aspects of each topic and how this changed with surrounding context and time.

Parallel coordinates (Inselberg, 2009) stack (interchangeable) vertical axes in parallel, each of which represents one attribute or variable (dimension) in a dataset. A polyline or polygon representing the attribute(s) of interest across all others intersects each axis at its value (cardinal or ordinal), providing a simple method for comparing trends between two adjacent or across large numbers of dimensions. While the interactive technique may be used to compare an infinite number of dimensions, practical limitations in screen width and resolution may require less important axes to be folded or hidden to allow a focus on more important attributes during detailed analysis of ROIs.

For each snapshot in the visualisation, the variable of interest is time, with six trend lines colour-coded (and annotated) in the same order across the parallel coordinates,

stacked with oldest (4th Aug) at the top. The key period is highlighted (in thick yellow): the 6th of Aug (the landing day of the Mars rover in 2012). The period immediately before each (5th) is highlighted in red, 4th Aug. is violet, 7th turquoise, 8th green and 9th blue. Because of the variation in tweet count between the two years, especially for Curiosity, we normalise each dataset during pre-processing, to aid comparison across the different batches. The second left axis is count or relative count where tweet counts are not uniform across the period (count for any day falls below the maximum set). The remaining vertical axes of the visualization represent one (simple or compound) term each.

The first round of the visual analysis used random samples of up to ten times (max. 2,000 per batch) those used for the content analysis. While the visualisation method used does not in itself impose an upper limit, this value was chosen to assess how representative of the overall dataset the trends observed in the smaller, manually annotated datasets were, and still have a sufficiently small dataset that allowed a good degree of manual inspection of its content. Further, this upper limit took into account the large variation in tweet count outside the immediate focus. Term selection was guided by knowledge of the content derived from the frames observed in the smaller, manually coded samples, and insight from alternative framing schemes. This was further refined using the larger sets of up to 2,000 each. Uren and Dadzie (2013) plot the entire dataset for a small set of terms for the Curiosity 2012 dataset, and Figure 2 shows plots for much larger samples, showing the results obtained by applying pattern matching across the larger datasets. (A larger set of snapshots comparing the trends can be found at <http://bit.ly/1DK3YkB>.)

Terms used to build the visualisation were extracted using pattern-matching. In addition to the knowledge of the data built during content analysis, patterns definition drew on knowledge of frames reported in the science communication literature. The benefit of the human in the loop is seen here; prior and contextual knowledge allows for more nuanced classification during manual inspection of ROIs. For instance, in Permafrost, a tweet comparing the revenue from oil exploitation in the Arctic to the environmental cost of melting permafrost at face value could appear to be “Economic”. However, following the frames described in (Nisbet and Scheufele, 2009), as it addresses controversial government plans for energy generation, there is a stronger argument for “Public accountability / governance”, than “Economic development/competitiveness”, so that the tweet is categorised as “Political”. Entradas *et al.* (2013), for instance, who discuss space exploration, provide additional frames we re-used especially in the analysis of the Curiosity dataset, to differentiate, for example, adventure from risk as the rover explores Mars.

Some terms use label and variant thereof, for example, for Curiosity “*Rover*” searches also for ‘*CuriosityRover*’. For *Phosphorus*, “*ELSI, Public accountability & Governance*” maps to $\{child|bomb|burn|war|crime|weapon|wmd|first\ aid\}$. To differentiate those from tweets also talking about white phosphorus as “*Political (public accountability)*” we look, for the latter, for where a specific target region or regime using the weapon is being discussed, of note being ‘ $\{fallujah|iraq|palestine|gaza|afghan|syria|pentagon|white\ house|regime\}$ ’ (or variants thereof, e.g., ‘palestine’ will also search for Palestinian).

Results

Figure 1 summarises the pertinent frames identified for each dataset and Table 2 presents the inter-coder agreement results. Note that the frame agreement values are more variable than batch agreement. For example, the *Fun* frame typically has only moderate agreement, possibly because humour is very personal (the coders had an eight year age difference and different cultural backgrounds), whereas the *Other languages* code had high agreement (note it is not 100% because some tweets use two languages, e.g., English and Arabic, and it was then down to the judgment of the coder whether they could assign it using the English text alone).

Dataset	Batch (SD)	Scientific	Political	Economic	ELSI	Fun	Off Topic	Other Languages
Cur2012	0.77(0.07)	0.72	0.50	0.00	0.20	0.58	0.73	0.98
Cur2013	0.83(0.05)	0.76	0.0	0.00	0.00	0.59	0.90	0.95
Phos2012	0.73(0.08)	0.74	0.66	0.70	0.76	0.51	0.74	0.85
Phos2013	0.90(0.10)	0.53	0.85	0.46	0.45	0.94	0.72	0.88
Perm2012	0.90(0.05)	0.77	1.00	0.33	0.89	0.60	0.95	0.95
Perm2013	0.71(0.06)	0.62	0.28	0.50	0.69	0.33	0.79	0.95
Frame(SD)		0.70(0.09)	0.55(0.34)	0.33(0.26)	0.50(0.31)	0.59(0.18)	0.80(0.09)	0.93(0.05)

Table 2 Inter-coder Agreement

The parallel coordinate plots are presented in pairs for each topic, with 2012 above 2013 (Figures 2-4). To allow direct comparison of relative trends for each attribute,

Figure 2: Visualization of the frequency of selected terms extracted from the Curiosity datasets, from the total number of tweets containing the filter ‘curiosity’, peaking at 67,703 on Aug. 6 2012 (upper plot), and 548 on the anniversary of the landing day in 2013 (lower plot). The middle plot filters out data in 2012 with frequency greater than 3500, revealing peaks outside Aug. 6 (which is subsequently greyed out). The lines are colour coded, one per day for each of Figures 2-4, and each is also annotated with a consistent icon to aid the reader in distinguishing the trend lines.

In Figure 2, the second tallest peak in 2013 represents tweets discussing the birthday theme with positive sentiment, using the complex, context-specific label *positive sentiment (Curiosity)*. This looks at the celebrations focusing on the birthday theme – *{year|anniversary|happy|birthday|song|sing}*. The absolute count in 2012 is still higher than that for the following year, picking up tweets using the word *happy* in the celebration on the landing day. However, as a percentage of total tweets, including the other specific terms classified as Fun, such as the *cat* jokes, this coordinate only shows a small rise in 2012. Notably, in 2013, while the term shows a consistent set of peaks for all days but the 4th, it quickly drops away after the anniversary. In 2012, ignoring the landing day, there is no noticeable difference in its usage across the period. Looking at the plots as a whole, the much larger number of tweets on the landing day in 2012 obscures other rises for the rest of the period, whereas much lower overall tweet count in 2013 makes it easier to distinguish variation in terminology usage across the period for 2013.

The trends in Figure 2 reflect our content analysis, with proportionally more tweeting

about the science behind the Mars rover and the landing event on its anniversary, and more celebration on the landing day. Figure 2 shows the dynamic evolution of eleven terms for the two periods out of a total of 29 terms for which a co-ordinate was generated in the process of data exploration. This is a ‘snapshot’ of the interactive visual analytic process. The terms selected represent ROIs, dimensions which, in this case, reveal areas of intense discussion.

The key difference between the two years is the marked increase in tweeting about the (trending) event on the landing day (in 2012), obscuring peaks for the surrounding days. The middle plot illustrates how, by lowering the upper limit of the plot and therefore greying out the dominant day, 06 Aug, the relatively smaller peaks for the surrounding days are more clearly revealed. 2013 sees much lower variation; there are however clear peaks the day prior to and on the anniversary itself, with a quick drop immediately after; publicity in the run up to the anniversary accounting for the brief renewal of public interest. There remains a good degree of sustained interest in the topic after the anniversary; as for 2012 this would be due to the attentive public and the scientific community. Explicit use of the word “science” (and variants thereof), for example, while dominant only on the landing day in 2012, show a clear set of peaks throughout the sample period in 2013, highest on the 5th and 6th.

Some terms were found to be strongly associated with either the original event only or its anniversary. Concern about the technology, and especially the landing gear, is captured in “*Technical Uncertainty*” (7th axis from the left in Figure 2); this peaked on the landing day but fell to near zero after the successful landing. ‘*Technical Innovation*’ (8th axis – concerning the camera and its output) sparked interest on the

landing day. Apart from a brief fall the following day, this was maintained as images of Mars were sent back to earth. In 2013, however, neither term recorded much interest. Some terms persist across both years, e.g., *Landing* (4th axis) and the Spanish word for Mars, *Marte* (5th axis), which with very high occurrence, indicates both international interest and the Spanish-speaking population in the US tweeting about the event.

Explicit “*congratulations*” shows a clear peak on 06 Aug 2012, but much smaller rises on the 5th and 6th in 2013. Terms for “*Fun*” popular in 2012, such as “*curiosity killed the cat*” and other cat jokes (12th axis) are almost non-existent in 2013. This prompted further analysis of the tweets framed as *Fun* in 2013; we found new jokes centred on the birthday song and theme. Tweets with ambiguous sentiment, for which the human element was necessary for disambiguation, include the terms used to collect additional information about ‘*Positive Sentiment*’ (9th axis) in 2013. Some of these referred to the scientific achievement / advanced technology that allowed the rover to sing happy birthday to itself, an event which some tweeters found poignant because it highlighted the lonely state of the rover, whose Twitter persona had an anthropomorphic tone. This prompted tweets such as those asking whether it had made friends with any aliens. Other terms which became more prominent in 2013 include *Adventure* (13th axis) – expressing the Rover’s year long exploration of the planet (while total numbers are smaller, the relative peak size rises to near 50% in 2013).

Phosphorus shows a swing from *Scientific* and *ELSI* frames dominating in 2012 to Fun and Political in 2013. The *Fun* tweets are predominantly periodic table jokes. This was not examined further because the scientific relevance was minimal, but we note it as a second example of the *Fun* frame linked to trending. Both the *Political* and the *ELSI* tweets are primarily associated with white phosphorus usage in Middle East conflicts. The difference between the years is that in 2012 the emphasis was on highlighting the effect of the weapons on civilians caught in the conflicts, whereas in 2013 a recent news item had introduced the terms “Pentagon” and “Syrian regime”. These mentions of political entities caused the tweets to be coded as *Political*. No single story dominated the tweets coded as *Scientific*, a wide range of subtopics was observed around agriculture, biology and school science (e.g., homework and exam revision).

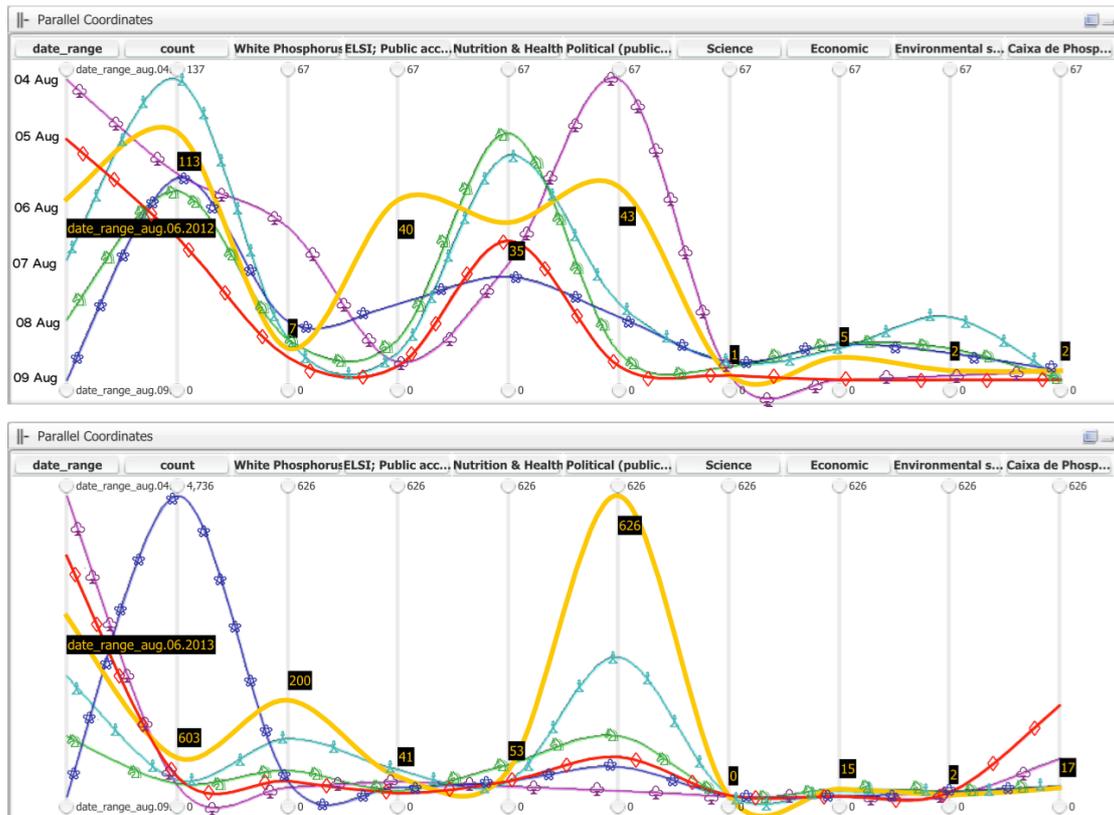


Figure 3: Frequently occurring terms in the Phosphorus datasets; the axis for jokes is hidden, as this dwarfs the other trend lines. (2012 upper plot,

2013 lower plot)

Filtering out off-topic spikes in the phosphorus 2013 dataset (the periodic table jokes) revealed smaller, relevant peaks. In Figure 3, the tallest set of peaks in 2013 are “*Political (public accountability)*” (6th axis from the left in Figure 3) and “*White Phosphorus*” (3rd axis); these related peaks concern contentious use of white phosphorus. The sudden jump to a peak on 6th Aug 2013 is due to retweeting of news that broke on the 5th, that the Pentagon (US) had admitted to using the weapon in Fallujah (Iraq). The focus on the regime and previous denial, over the moral implications of the action, is the key to framing as Political (concerning public accountability) rather than ELSI. While the use of the term continues to fall for the rest of the period it is the only consistent set of peaks in 2013.

Looking at 2012, the smaller peaks for other terms are more apparent, not being dwarfed by a single spike as in 2013. Starting on 4 Aug the highest peak is also “*Political (public accountability)*”; this is again due largely to “*White Phosphorus*”. The latter however stays relatively low for the rest of the period. The other set of peaks of note is “*Nutrition & Health*” (5th axis), framed as *ELSI* as they fall under social issues or popular science/psychology. Such tweets discuss or advise on homeopathy, health and nutrition, e.g. in lists of nutrients and minerals found in particular foods. “*Nutrition & Health*”, although not the highest, comprises the only consistent set of peaks throughout 2012. This, with a small contribution from “*Environmental stewardship*” (9th axis), accounts for the relatively high percentage of tweets framed as *ELSI* in 2012 (see Figure 1).

Permafrost had a large number of *Off Topic* tweets in 2012. Overall, numbers of tweets were low, and hence one set of retweets linking to a social documentary entitled “*Conquering Permafrost: People of the BAM railway*” became prominent. More interesting (for our purposes) frames were *Scientific* and *ELSI*, which were well populated in both the Permafrost datasets. The *Scientific* tweets predominantly reported results from scientific studies. Typically, these reported record melting, in particular following an unusually hot summer in 2013. The *ELSI* tweets were broadly of two types, the first concerned the effects of permafrost melt on communities in the high north, the second type point to the moral responsibility to take action on climate change.

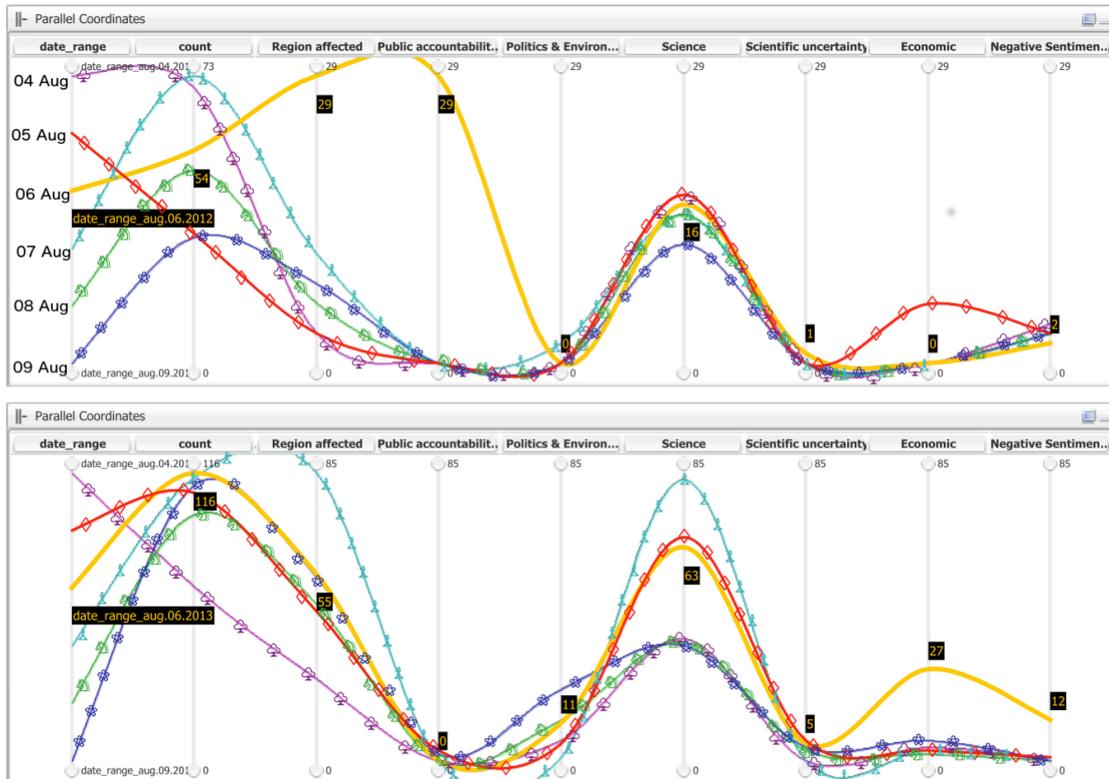


Figure 4: Frequently occurring terms within the Permafrost datasets (2012 upper plot, 2013 lower plot)

In Figure 4, the visualisation detects an ROI for 06 Aug. 2012: two significant peaks are seen for terms to do with “*Public accountability, Environmental stewardship & Ethics*” (4th axis from the left in Figure 4) and “*Region affected*” (3rd axis – including countries or regions frequently mentioned). All trend lines for both years show peaks for “*Science*” (6th axis), which detects terms related to permafrost melt. The steepest peaks are for 7 Aug. 2013, focusing on the effect on the environment and the residents of the geographical areas concerned. Finally, “*Economic*” (8th axis) sees small peaks on 5 Aug. 2012 and 6 Aug. 2013 but is otherwise not well populated. That for 5 Aug. 2012 is interesting in that it maps to retweets about the controversy surrounding the proposed Enbridge gas pipeline in Canada, due to potential socio-cultural and climate risks.

Limitations

This study has considered only three scientific topics sampled in two consecutive years. Although these illustrate a range of frames they cannot be representative of all scientific topics, having been selected only to illustrate trending and non-trending topics. It is not possible to conclude with absolute certainty, from the data gathered, what the lower sample size limits are.

Parallel coordinates are often found to have a relatively high learning curve. Further, with superficial visual similarity to line graphs, static snapshots in print communications can make interpretation non-intuitive. The benefits in the approach are better seen in actual use, as the interactive components allow trends and ROIs to be recognised as the dynamic filters are applied and co-ordinates relocated to allow

focus on and detailed analysis of selected attributes. We acknowledge that until the visualization becomes more familiar, the reader may require more detailed guidance for interpretation.

Finally, in the approach reported here, the dimensions/attributes analysed were generated using hand-coded patterns, albeit guided by research on framing and the visual overviews. Augmenting this with automatically generated dimensions, for example, produced by co-word clustering, is the obvious next step. More advanced social media metrics could also be deployed. For example, in previous studies we have experimented with the use of ageing factors, calculated using retweets, to study attention to science events in the news (Uren and Dadzie, 2012).

Conclusions

The first research question was whether this method supports the analysis of dynamic changes in non-trending topics. That Twitter has potential to monitor engagement with newsworthy science topics or big science events, like Curiosity, is expected. Studies of science communication suggest that space exploration events typically generate significant interest (Baram-Tsabari and Segev, 2011), and that the sector of the public that is interested in science has a positive attitude, overall, toward space exploration (Entradas *et al.*, 2013). We have, further, demonstrated that we can observe dynamic changes to the framing of science communication for less populous topics, such as Phosphorus and Permafrost, which had only tens of tweets on some of the days sampled.

The second research question concerned whether changes could be observed across disconnected time periods. We found that a longitudinal view brought out the development of the debates. The combination of multi-dimensional visualisation and pattern matching allowed rapid discovery and analysis of communication patterns – term emergence and usage – within each frame and time period. For Curiosity we saw how the audience shrank to an attentive group, with more interest in science issues, but still retaining enthusiasm for the adventure of scientific exploration. For Phosphorus and Permafrost we saw how political frames grew on Twitter and how polarised they can be, with particular viewpoints and arguments dominant in both topics. For example, only one side of the climate change debate was much represented. Hestres (2013) has discussed the tactics of climate advocacy campaigns, and our samples suggest they have been successful in engaging an “issue public” around the negative impact of permafrost melt on affected regions.

The third question asked whether visualisation-based analysis reveals further information in addition to confirming the content analysis. We found that the multi-dimensional visualisation technique, parallel coordinates, used with compound terms, provides a candidate method to support content analysis approaches to the interpretation of frames. This exploratory process allowed fluid addition of new terms to the analysis. For example, in the Curiosity dataset we identified a shift from issues of technical uncertainty and the landing event itself in 2012, to the use of terms around scientific exploration as an adventure in 2013. The higher degree of automation also allowed us to analyse bigger samples in detail (up to 2,000 tweets per batch, compared to 200 in the content analysis), and to examine smaller time slices

(day by day in this report as opposed to the two six day periods). From a practical viewpoint, the visualisation-based approach reduces the labour required for pure content analysis while increasing breadth and coverage.

This work is positioned at the intersection of science communication studies and social media metrics. The research questions sought to analyse the suitability of the parallel coordinates visualization, in combination with pattern matching, for the analysis of the framing of microblogs – a science communication activity. The analyses presented here indicate that it is an effective method for observing dynamic changes in communication on Twitter, with the scalability required for longitudinal studies, and the flexibility to study medium scale datasets. The visualization approach has potential for use with a range of social media metrics with a view to improving the understanding of scientific contents discussed on Twitter. Exploring these further is the next stage of our research.

References

- Adams, A., Lomax, G. and Santarini, A. (2011), “Social media & stem cell science: examining the discourse”, *Regenerative Medicine*, Vol. 6 No. 6s, pp. 121–124.
- Bar-Ilan, J., Haustein, S., Peters, I., Priem, J., Shema, H. and Terliesner, J. (2012). Beyond citations: Scholars’ visibility on the social web”, *Proceedings of the 17th International Conference on Science and Technology Indicators*, Repro-UQAM, Montreal, Canada, pp. 98–109.
- Baram-Tsabari, A. and Segev, E. (2011), “Exploring new web-based tools to identify public interest in science”, *Public Understanding of Science*, Vol. 20 No. 1, pp. 130–145.
- Bauer, M.A., Allum, N. and Miller, S. (2007), “What can we learn from 25 years of PUS survey research? Liberating and expanding the agenda”, *Public Understanding of Science*, Vol. 16 No. 1, pp. 79–95.
- Börner, K., Chen, C. and Boyack, K.W. (2003), “Visualizing knowledge domains”, *Annual Review of Information Science and Technology*, Vol. 37 No. 1, pp. 179–255.

- Bruu Carver, R., Rodland, A. and Breivik, J. (2013), “Quantitative frame analysis of how the gene concept is presented in tabloid and elite newspapers”, *Science Communication*, Vol. 35 No. 4, pp. 449–475.
- Cano, A.E., Dadzie, A.-S. and Ciravegna, F. (2013), “Travel mashups”, in Endres-Niggemeyer, B. (Ed), *Semantic Mashups. Intelligent Reuse of Web Resources*, Springer Verlag, Berlin/Heidelberg, pp. 321–347.
- Cao, N., Lin, Y.-R., Sun, X., Lazer, D., Liu, S. and Qu, H. (2012), “Whisper: Tracing the spatiotemporal process of information diffusion in real time”, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 18 No. 12, pp. 2649–2658.
- Chew, C. and Eysenbach, G. (2010), “Pandemics in the age of Twitter: Content analysis of tweets during the 2009 H1N1 outbreak”, *PLoS ONE*, Vol. 5 No. 11, e14118.
- Chua, F.C.T. and Asur, S. (2013), “Automatic summarization of events from social media”, in *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, AAAI Press, Palo Alto, CA, pp. 81–90.
- Cobo, M.J., López-Herrera, A.G., Herrera-Viedma, E. and Herrera, F. (2012), “SciMat: a new science mapping analysis software tool”, *Journal of the American Society for Information Science and Technology*, Vol. 63 No. 6, pp. 1609–1630.
- Cohen, J. (1960), “A coefficient of agreement for nominal scales”, *Educational and Psychological Measurement*, Vol. 20 No. 1, pp. 37–46.
- Cui, W., Liu, S., Tan, L., Shi, C., Song, Y., Gao, Z., Qu, H. and Tong, X. (2011), “TextFlow: Towards better understanding of evolving topics in text”, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 17 No. 12, pp. 2412–2421.
- Darling, E.S., Shiffman, D., Côté, I.M. and Drew, J.A. (2013), “The role of Twitter in the life cycle of a scientific publication”, *Ideas in Ecology and Evolution*, Vol. 6 No. 1, pp. 32–43.
- Desai, T., Shariff, A., Shariff, A., Kats, M., Fang, X., Christiano, C. and Ferris, M. (2012), “Tweeting the meeting: An in-depth analysis of Twitter activity at Kidney Week 2011”, *PLoS ONE*, Vol. 7 No. 7, e40253.
- de Domenico, M., Lima, A., Mougél, P. and Musolesi, M. (2013), “The anatomy of a scientific rumor”, *Scientific Reports*, Vol. 3, 2980.
- Dou, W., Yu, L., Wang, X., Ma, Z. and Ribarsky, W. (2013), “HierarchicalTopics: visually exploring large text collections using topic hierarchies”, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 19 No. 12, pp. 2002–2011.
- Entradas, M., Miller, S. and Peters, H.P. (2013), “Preaching to the converted? An analysis of the UK public for space exploration”, *Public Understanding of Science*, Vol. 22 No. 3, pp. 269–286.
- Eysenbach, G. (2011), “Can tweets predict citations? Metrics of social impact based on Twitter and correlation with traditional metrics of scientific impact”, *Journal of Medical Internet Research*, Vol. 13 No. 4, e123.

- Fahy, D. and Nisbet, M.C. (2011), "The science journalist online: Shifting roles and emerging practices", *Journalism*, Vol. 12 No. 7, pp. 778–793.
- Goffman, E. (1974), *Frame Analysis: An Essay on the Organization of Experience*, Harper and Row, London.
- Haustein, S., Bowman, T.D., Holmberg, K., Peters, I. and Larivière, V. (2014a), "Astrophysicists on Twitter: an in-depth analysis of tweeting and scientific publication behaviour", *Aslib Journal of Information Management*, Vol. 66 No. 3, pp. 279–296.
- Haustein, S., Peters, I., Sugimoto, C.R., Thelwall, M. and Larivière, V. (2014b), "Tweeting biomedicine: An analysis of tweets and citations in the biomedical literature", *Journal of the Association for Information Science and Technology*, Vol. 65 No. 4, pp. 656–669.
- Hermida, A. (2010), "Twittering the news. The emergence of ambient journalism", *Journalism Practice*, Vol. 4 No. 3, pp. 297–308.
- Hestres, L.E. (2013), "Preaching to the choir: Internet-mediated advocacy, issue public mobilization, and climate change", *New Media Society*, Vol. 16 No. 2, pp. 323-339.
- Hogan, N.M. and Sweeney, K.J. (2013), "Social networking and scientific communication: A paradoxical return to Mertonian roots?", *Journal of the American Society for Information Science and Technology*, Vol. 64 No. 3, pp. 644–646.
- Hooper, R.S. (1965), *Indexer Consistency Tests - Origin, Measurements, Results and Utilization*, IBM Corp, Bethesda, MD.
- Hu, Y., John, A., Seligmann, D.D. and Wang, F. (2012), "What were the tweets about? Topical associations between public events and Twitter feeds", in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, AAAI Press, Palo Alto, CA, pp. 154–161.
- Hu, Y., Talamadupula, K. and Kambhampati, S. (2013), « Dude, srsly?: The surprisingly formal nature of Twitter's language », in *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, AAAI Press, Palo Alto, CA, pp. 244–253.
- Hubmann-Haidvogel, A., Brasoveanu, A.M.P., Scharl, A., Sabou, M. and Gindl, S. (2012), "Visualizing contextual and dynamic features of micropost streams", in *Proceedings of the WWW'12 Workshop on Making Sense of Microposts (MSM2012)*, Lyon, pp. 34–40.
- Inselberg, A. (2009), *Parallel Coordinates: Visual Multidimensional Geometry and Its Applications*, Springer-Verlag, New York.
- Kraker, P., Weißensteiner, P. and Brusilowsky, P. (2014), "Altmetrics-based visualizations depicting the evolution of a knowledge domain", in *Proceedings of the 19th International Conference on Science and Technology Indicators (STI 2014)*, Universiteit Leiden, The Netherlands, pp. 330-333.
- Krippendorff, K. (2004), *Content Analysis: An Introduction to its Methodology*, Sage, Thousand Oaks, CA.
- Leininger, K. (2000), "Interindexer consistency in PsycINFO", *Journal of Librarianship and Information Science*, Vol. 32 No. 1, pp. 4–8.

- Letierce, J., Passant, A., Decker, S. and Breslin, J.G. (2010), “Understanding how Twitter is used to spread scientific messages”, in *Proceedings of the Web Science Conference (WebSci'10) Extending the Frontiers of Society On-Line*, Raleigh, NC.
- Lumezanu, C., Feamster, N. and Klein, H. (2012), “#bias: Measuring the tweeting behavior of propagandists”, in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, AAAI Press, Palo Alto, CA, pp. 210–217.
- Mandavilli, A. (2011), “Trial by Twitter”, *Nature*, Vol. 469 No. 7330, 286–287.
- Medelyan, O. and Witten, I.H. (2006), “Measuring inter-indexer consistency using a thesaurus”, in *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries*, Chapel Hill, NC, pp. 296–297.
- Meyer, E. and Schroeder, R. (2009), “The world wide web of research and access to knowledge”, *Knowledge Management Research and Practice*, Vol. 7 No. 3, pp. 218–233.
- Miller, J.D. (1983), *The American People and Science Policy: The Role of Public Attitudes in the Policy Process*. Pergamon Press, New York, NY.
- Neset, T.-S.S. and Cordell, D. (2012), “Global phosphorus scarcity: identifying synergies for a sustainable future”, *Journal of the Science of Food and Agriculture*, Vol. 92 No. 1, pp. 2–6.
- Nisbet, M.C. (2009), “Communicating climate change: why frames matter for public engagement”, *Environment*, Vol. 51 No. 2, pp. 12–23.
- Nisbet, M.C. and Scheufele, D.A. (2009), “What’s next for science communication? Promising directions and lingering distractions”, *American Journal of Botany*, Vol. 96 No. 10, pp. 1767–1778.
- Ponte, D. and Simon, J. (2011), “Scholarly communication 2.0: Exploring researchers’ opinions on web 2.0 for scientific knowledge creation, evaluation and dissemination”, *Serials Review*, Vol. 37 No. 3, pp. 149–156.
- POST (Parliamentary Office of Science and Technology) (2003), “Science and society: Three years on”, available at: <http://www.parliament.uk/documents/post/13-may-proceedings.pdf> (accessed 8 March 2015).
- Priem, J. and Hemminger, B.M. (2010), “Scientometrics 2.0: toward new metrics of scholarly impact on the semantic web”, *First Monday*, Vol. 15 No. 7, available at: <http://firstmonday.org/article/view/2874/2570> (accessed 13 March 2015).
- Priem, J., Piwowar, H.A. and Hemminger, B.M. (2012), “Altmetrics in the wild: Using social media to explore scholarly impact”, available at: <http://arxiv.org/abs/1203.4745> (accessed 8 March 2015).
- Reddy Yerva, S., Miklós, Z. and Aberer, K. (2011), “What have fruits to do with technology? The case of Orange, Blackberry and Apple”, in *International Conference on Web Intelligence, Mining and Semantics, WIMS'11*, Sogndal, Norway, 48.
- Ritter, A., Clark, S., Mausam and Etzioni, O. (2011), “Named entity recognition in tweets: An experimental study”, in *Proceedings of the Conference on Empirical*

- Methods in Natural Language Processing*, Association for Computational Linguistics, Stroudsborg, PA, pp. 1524-1534.
- Roemer, C.R. and Borchart, R. (2013), "Institutional altmetrics and academic libraries", *Information Standards Quarterly*, Vol. 25 No. 2, pp. 14–19.
- Schaefer, K., Lantuit, H., Romanovsky, V. and Schuur, E.A.G. (2012), *Policy Implications of Warming Permafrost, United Nations Environment Programme Special Report*, Nairobi, Kenya.
- Schäfer, M.S. (2009), "From public understanding to public engagement: An empirical assessment of changes in science coverage", *Science Communication*, Vol. 30 No. 4, 475-505.
- Shuai, X., Pepe, A. and Bollen, J. (2012), "How the scientific community reacts to newly submitted preprints: Article downloads, Twitter mentions, and citations", *PLoS ONE*, Vol. 7 No. 11, e47523.
- Thelwall, M., Buckley, K. and Paltoglou, G. (2011), "Sentiment in Twitter events", *Journal of the American Society for Information Science and Technology*, Vol. 62 No. 2, pp. 406–418.
- Thelwall, M., Haustein, S., Larivière, V. and Sugimoto, C.R. (2013), "Do altmetrics work? Twitter and ten other social web services", *PLoS ONE*, Vol. 8 No. 5, e64841.
- Tufekci, Z. (2014), "Big questions for social media big data: Representativeness, validity and other methodological pitfalls, in *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, AAAI Press, Palo Alto, CA, pp. 505–514.
- Uren, V. and Dadzie, A.-S. (2011), "Relative trends in scientific terms on Twitter", in *altmetrics11: Tracking scholarly interest on the Web, Workshop at ACM WebSci'11*.
- Uren, V. and Dadzie, A.-S. (2012), "Ageing factor: a potential altmetric for observing events and attention spans in social media", in *Proceedings of the Knowledge Extraction and Consolidation from Social Media Workshop at ISWC 2012*.
- Uren, V. and Dadzie, A.-S. (2013), "Nerding out on Twitter: Fun, patriotism and #Curiosity", in *Proceedings of the 22nd international conference on World Wide Web*, pp. 605-612.
- Van Eck, N. J. and Waltman, L. (2010), "Software survey: VOSviewer, a computer program for bibliometric mapping", *Scientometrics*, Vol. 84 No. 2, pp. 523–538.
- Van Eck, N. J. and Waltman, L. (2014), "Visualizing bibliometric networks", in Ding, Y., Rousseau, R. and Wolfram, D. (Eds.), *Measuring Scholarly Impact: Methods and Practice*, Springer, Berlin, pp. 285–320.
- Veltri, G. A. (2013), "Microblogging and nanotweets: Nanotechnology on Twitter", *Public Understanding of Science*, Vol. 22 No. 7, 832–849.
- Weller, K., Dröge, E. and Puschmann, C. (2011), "Citation analysis in Twitter. Approaches for defining and measuring information flows within tweets during scientific conferences", in *Proceedings of the ESWC'11 Workshop on Making Sense of Microposts (MSM2011)*, Lyon, 1–12.

- Wilkinson, D. and Thelwall, M. (2012), “Trending Twitter topics in English: An international comparison”, *Journal of the American Society for Information Science and Technology*, Vol. 68 No. 8, pp. 1631–1646.
- Xu, P., Wu, Y., Wei, E., Peng, T., Liu, S., Zhu, J. J. H. and Qu, H. (2013), “Visual Analysis of Topic Competition on Social Media“, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 19 No. 12, pp. 2012–2021.
- Yuan, X., Guo, P., Xiao, H., Zhou, H. and Qu, H. (2009), “Scattering points in parallel coordinates”, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 15 No. 6, pp. 1001–1008.
- Zahedi, Z., Costas, R. and Wouters, P. (2014), “How well developed are altmetrics? A cross-disciplinary analysis of the presence of “alternative metrics” in scientific publications”, *Scientometrics*, Vol. 101 No. 2, pp. 1491–1513.