

Quantitative Approaches to Comparing Communication Patterns on Twitter

Dr. Axel Bruns

Associate Professor

ARC Centre of Excellence for Creative Industries and Innovation

Queensland University of Technology

Brisbane, Australia

a.bruns@qut.edu.au – @snurb_dot_info – <http://mappingonlinepublics.net/>

Phone: +61 7 3138 5548

Dr. Stefan Stieglitz

Assistant Professor

Department of Information Systems

University of Münster

Münster, Germany

stefan.stieglitz@uni-muenster.de – <http://www.wi-kuk.de/>

Phone: +49 251 83 38 115

Abstract

To date, available literature mainly discusses Twitter activity patterns in the context of individual case studies, while comparative research on a large number of communicative events, their dynamics and patterns is missing. By conducting a comparative study of more than forty different cases (covering topics such as elections, natural disasters, corporate crises, and televised events) we identify a number of distinct types of discussion which can be observed on Twitter. Drawing on a range of communicative metrics, we show that thematic and contextual factors influence the usage of different communicative tools available to Twitter users, such as original tweets, @replies, retweets, and URLs. Based on this first analysis of the overall metrics of Twitter discussions, we also demonstrate stable patterns in the use of Twitter in the context of major topics and events.

Keywords

Twitter, social media, public communication, media events, communicative patterns

Quantitative Approaches to Comparing Communication Patterns on *Twitter*

Introduction

Since 2006, microblogging has become an increasingly widely used tool for communication on the Internet. *Twitter*, as one of the first and most popular microblogging providers, has some 140 million users, with some 340 million tweets posted each day (*Twitter*, 2012). In contrast to social networking sites (SNS) such as *Facebook*, the reach of posts on *Twitter* is not necessarily limited to a specific group (such as subscribed 'friends' or 'followers'); rather, posted messages are public by default and may also be found by visitors searching the site or tracking the *Twitter* stream. Each user is thus able to create public posts to initiate discussions, to participate in debates, and to follow the communication of others. To manage these communicative flows and increase the efficiency of public message exchanges, *Twitter* users have adapted a variety of methods to classify their contributions (tweets) – for example as a public response (@reply) or a shared message originating from another user (retweet).

Twitter has now become a widely used communications channel across a wide range of applications, from politics, journalism and crisis communication (e.g. Bruns & Burgess, 2011a; Christensen, 2011; Larsson & Moe, 2011; Lotan *et al.*, 2011; Bruns *et al.*, 2012; Mendoza *et al.*, 2010; Palen *et al.*, 2010; Stieglitz & Dang-Xuan 2012) through its use as a backchannel for television shows, cultural and sporting events, and conferences to a wide variety of uses for everyday interpersonal communication (e.g. Deller, 2011; Dröge *et al.*, 2011; Weller *et al.*, 2011; boyd *et al.*, 2009; Papacharissi, 2011; Marwick & boyd, 2011). Significant research into some such uses is now emerging, but largely remains in the form of topic-, context-, and event-related case studies which are able to shed substantial light on specific uses of *Twitter*, but do not yet lead to a more comprehensive overall picture of how *Twitter* is used.

Individuals and organizations may use *Twitter* to subscribe to the update feeds of other users, as well as to publish their own short messages (to a maximum of 140 characters) about various topics (e.g. from personal and professional updates to press releases and other corporate information). To widely disseminate information on *Twitter*, the mechanism of retweeting has been adopted by users. By retweeting, users may

not only share information but also entertain a certain audience or (by adding comments to the retweets) publicly agree or disagree with someone (boyd *et al.*, 2010). As a result, *Twitter* has become an important platform for users to spread information about topics of shared interest: retweets propagate the original tweet to a new set of audiences, namely the followers of the retweeting user. Given the growing *Twitter* userbase, the high speed of information dissemination on *Twitter*, and the significant influence of *Twitter* as a driver of Web traffic, new questions arise about the way it is used to support public information sharing and information search, therefore.

Other studies have already made some first steps to investigate how and why certain information items spread more widely than others (Suh *et al.*, 2010; Stieglitz & Dang-Xuan, 2012). However, so far the literature mainly discusses *Twitter* activity patterns in the context of individual case studies, while comparative research on a large number of discussions, their dynamics and patterns is missing. By conducting a comparative study of several dozen different cases (including topics such as elections, natural disasters, corporate crises, and televised sporting and cultural events) we have identified a number of distinct types of discussion which can be observed on *Twitter*. Drawing on a range of communicative metrics, we show that thematic and contextual factors influence the usage of different communicative tools available to *Twitter* users, such as original tweets, @replies, retweets, and URLs. Furthermore, we also demonstrate patterns in the structure of the user community involved (e.g. number of participants, relevance of lead users). This paper presents a first analysis of the overall metrics of *Twitter* discussions relating to different areas of content, and outlines two standard, stable types of using *Twitter* in the context of major topics and events. As such, it represents a significant advance for research which investigates the usage of different communication tools in public discussions.

This article pursues this larger picture by exploring general patterns of *Twitter* usage, drawing on detailed usage metrics for a wide range of cases and events over the past two years. By collating these data points and identifying cases which exhibit similar patterns of activity, we observe a range of common, apparently well-established user practices on *Twitter*. We suggest that these observations point to regularities in the popular responses to specific themes and events that may also be identified well beyond the *Twitter* platform itself.

Related Work

In recent years, a substantial amount of literature has been published in the field of *Twitter* communication. Therefore, we provide a short literature review of those articles which explicitly reflect metrics within *Twitter* communication in the field of politics, natural and human disasters, as well as entertainment and brand-related communication.

Politics

In a study of approximately 100,000 messages containing a reference to either a political party or a politician in the context of the 2009 German federal election, Tumasjan *et al.* (2011) show that *Twitter* is used extensively for the dissemination of politically relevant information and that the mere number of party mentions accurately reflects the election result, suggesting that microblogging messages on *Twitter* seem to validly mirror the political landscape offline and can be used to predict election results to a certain extent. Conover *et al.* (2011) examine two networks of political communication on *Twitter* with more than 250,000 tweets from the six weeks leading up to the 2010 U.S. congressional midterm elections. Using a combination of network clustering algorithms and manually-annotated data, the authors demonstrate that the network of political retweets exhibits a highly segregated partisan structure, with extremely limited connectivity between left- and right-leaning users. Surprisingly, this is not the case for the user-to-user mention network, which is dominated by a single politically heterogeneous cluster of users in which ideologically-opposed individuals interact at a much higher rate compared to the network of retweets. Similarly, Yardi and boyd (2010) find that, in a political context, *Twitter* users are more likely to interact with others who share the same views as they do in terms of retweeting, but they are also actively engaged with those with whom they disagree. In addition, replies between like-minded individuals would strengthen group identity, whereas replies between different-minded individuals would reinforce in-group and out-group affiliation. In a large-scale study, Suh *et al.* (2010) addressed these questions and identified several factors that significantly impact on the retweetability of *Twitter* messages (tweets), including the presence of URLs and hashtags, as well as the number of followers and the age of the originating user's account. Beside these case-based analyses, Stieglitz & Dang-Xuan (forthcoming) provide a general framework which presents methods for social media analytics in the political context.

Natural and Human Disasters

In recent years, a growing body of literature has emerged in the field of social media and crisis communication (Bruns *et al.*, 2012; Hughes & Palen, 2009; Mendoza *et al.*, 2010; Palen *et al.*, 2010). Cheng *et al.* (2011) investigated *Twitter* as a tool to monitor and capture emerging trends and patterns of time-critical knowledge. They extensively evaluated a diffusion-based recommendation framework and a proposed algorithm using *Twitter* data collected during the early outbreak of H1N1 Flu. Studies by Bruns *et al.* (2012) and Cheong & Cheong (2011) analysed *Twitter*-based communication in the context of natural disasters, focussing on the Australian floods in 2011; by using social network analysis methods they found that several different groups of actors, including affected locals, emergency services, and mainstream media organisations, played important roles in providing and sharing information about the disaster.

Other studies (Vis, 2012; Lotan *et al.*, 2011; Bruns *et al.*, forthcoming) examined uses of *Twitter* during major civil unrest in the context of the 2011 London and UK riots or the Arab Spring uprisings in a number of North African and Middle East countries. They identified a diverse range of uses of social media for information dissemination alongside and in addition to other media channels and word-of-mouth information, and also highlighted significant differences in activity between local and more distant observers of these events.

Entertainment and Brand-Related Communication

Stieglitz and Krüger (2011) investigated a 2010 brand crisis involving car manufacturer Toyota and showed that, measured by the published number of tweets, crisis discussions are characterised by peaks and quiet periods in the communication of enterprise-related issues. Further, they found that the lead users involved in the *Twitter* debate played an important role in the discussion (e.g. by publishing a significantly high amount of all tweets and generating a large amount of retweets).

Park *et al.* (2009) investigated the Domino's Pizza crisis and analysed the diffusion of bad news through *Twitter*. They separately classified sentiments in tweets generated by customers and those generated by the enterprise itself. Based on their study they prove that the diffusion of bad news is faster than that of other types of content (e.g. apologies).

In a study of users and their behaviours in the *Twitter* network, Krishnamurthy *et al.* (2008) identified three types of users (broadcaster, acquaintances and miscreants) by analysing a crawled data set that covered nearly 100,000 users. The broadcasters, also named power-tweeters, are characterised by a large number of followers as well as a large amount of self-created postings. One finding in this study was that these users update their status more often and post more tweets than users of the two other categories.

Methodology

In order to establish a sound basis for the identification of shared patterns in *Twitter*-based communication around specific issues, the majority of the *Twitter* phenomena which we observe in this article are centred around common hashtags (brief keywords included in tweets, prefixed with the hash symbol '#'). Hashtags are an originally user-generated mechanism for making messages related to a specific topic more easily discoverable, and are now well-supported by central *Twitter* infrastructure as well as by specific *Twitter* client software; it is now possible for users (and even for non-registered visitors to the site) to search *Twitter* for specific hashtags, and to follow the stream of new messages containing specific hashtags in real time. This makes hashtags a useful and an important mechanism for coordinating conversations around identified themes and events, ranging from breaking news (such as #eqnz for the 2010/11 earthquakes in Christchurch, New Zealand) through major media events (e.g. #euro2012 for the 2010 European Football Championships) to viral marketing campaigns (such as #kony2012 and #stopkony for the campaign to bring a fugitive Ugandan warlord to justice). Beyond such world events, hashtags are also used to coordinate much more low-key discussions and user communities, from providing a backchannel for conference delegates to organising *Twitter*-based user meetups (such as the long-standing #phdchat, a global discussion for PhD candidates). Finally, a different use of hashtags – which we do not consider in detail here – is as markers of emphasis or emotion (as in, “My bus is running late again. #fail”).

Hashtags, then, may emerge *ad hoc* in response to breaking news and other unforeseen events, spreading virally as more and more users with an interest in the topic see the hashtag in their *Twitter* feeds and begin to use it themselves (see Bruns & Burgess, 2011b); they may also be used repeatedly for recurring events (such as #ausvotes for Australian federal elections, or #eqnz for each of the four major earthquakes which affected Christchurch in 2010/11); or they may be promoted *praeter hoc* by relevant organisations as the appropriate

hashtag to be used for an upcoming event (this is the case for backchannel hashtags for conferences or TV shows, for example). Such diverse hashtags may in turn attract widely varying groups of users: breaking news events, especially where they are of national or global relevance, may find hundreds of thousands of *Twitter* users posting or retweeting hashtagged messages, while hashtags related to conferences or TV shows may involve only a much smaller number of users who happen to be attending or watching at the time. Standing hashtags for the discussion of specific continuing topics (from #phdchat to the day-to-day tracking of long-term events such as the popular revolts in #Libya, #Egypt or #Syria), in turn, may involve only a comparatively small group of committed contributors, but see a temporary influx of a large number of interested users as key developments unfold and are widely covered by mainstream media outlets.

Using the *Twitter* Application Programming Interface (API), it is comparatively simple to capture comprehensive datasets of the vast majority of all tweets containing a specific hashtag (within limits determined by the reliability of the API and of real-time *Twitter* tracking tools; cf. Bruns & Liang, 2012); during 2010-12, as part of a project collaboration between Queensland University of Technology, Brisbane, and the University of Münster, we have done so for some 40 hashtags as well as a number of non-hashtagged keywords (which we will discuss later). Individual hashtag datasets for this study were captured using the open-source platform *yourTwrapperkeeper* (2012), which utilises *Twitter* streaming API and search API functionality to capture, in real time, any tweets containing the keywords (including hashtags) selected by the operator (see Bruns & Liang, 2012, for more details on *Twitter* research methods using *yourTwrapperkeeper* and alternative technologies). *yourTwrapperkeeper* does not provide for *post hoc* data gathering; it is able only to capture tweets for set keywords as they are sent, and the selection of hashtag and keyword datasets used for this study is a function of the long-term research interests of the Brisbane and Münster research groups, therefore, which specialise in political, crisis, and brand communication research. In themselves, these over 40 datasets cover a diverse range of uses, therefore, but we also encourage the further extension of this initial work through the addition of hashtag and keyword metrics extracted from datasets gathered by researchers interested in other areas of communication using social media.

To better understand the diversity of uses evident in the present collection of cases, and to identify any common patterns between individual cases, we draw on a catalogue of metrics for describing the communicative patterns which may be observed for each hashtag (see Bruns & Stieglitz, forthcoming, for a detailed introduction of these metrics). In the first place, these include:

- the number of *tweets* in the hashtag dataset;
- the number of *unique users* contributing to the hashtag dataset;
- the percentage of *original tweets* in the hashtag dataset (i.e. tweets which are neither @replies nor retweets);
- the percentage of *genuine @replies* in the hashtag dataset (i.e. @replies which are not retweets);
- the percentage of *retweets* in the hashtag dataset¹; and
- the percentage of tweets in the hashtag dataset which contain *URLs*.

Additionally, we also divided the total userbase for each hashtag dataset into three groups, following a standard 1/9/90 distribution (Tedjamulia *et al.* 2005):

- the top 1% of most active *lead users*;
- the next 9% of still *highly active users*;
- and the remaining 90% of *least active users*.

For each of these three groups in each hashtag dataset, we again calculated the metrics outlined above, taking into account only the tweets sent by that percentile group. Compared across the groups, this provides a measure for each hashtag of how dominant within the overall hashtag conversation the leading user groups are, and enables us to examine any obvious differences in the Twitter activity patterns of the three user groups.

In the following discussion, we collate and compare these metrics for the range of hashtags which we tracked over the past two years. This enables us to identify communication patterns which are common across these diverse cases, and to develop a typology of hashtagged *Twitter* usage. First, however, we provide an overview of the hashtag datasets which were used in this analysis, and outline their relevant features.

¹ A number of equivalent user conventions for marking messages as retweets now exist, and are included in this figure: in addition to the most common variant RT @user [message], we tested for MT @user [message] (for 'manual tweet'), [message] (via @user), and "@user [message]" (username and original message enclosed in quotation marks) are also common, with or without added comments from the retweeting user. A second retweeting mechanism was introduced by *Twitter* itself, by providing a 'retweet button' next to each tweet displayed on its Website or in *Twitter* clients; this mechanism passes along the original tweet verbatim and in full length, without inserting the "RT @user" into the message. Many *Twitter* clients – including the version of the *Twitter* Website optimised for mobile devices – now offer a choice between both formats (cf. Bruns, 2012). Because, contrary to 'manual' retweets, such 'button' retweets do not result in a new message, but simply add to the metadata of the original tweet, they are not captured by our *Twitter* tracking solution, and are therefore excluded from the datasets it captures. This is an unavoidable gap in the datasets, resulting in a systematic underestimation of retweeting activity. It is possible, however, to extrapolate overall retweeting activities from the patterns of 'manual' retweeting.

Hashtag Datasets

This study draws on a wide variety of datasets, whose key features we outline in Table 1:

[insert table 1 here]

While the scope of this article does not permit a detailed discussion of the themes and contents of each dataset, we note that these cases encompass a wide variety of topical hashtag uses – they range from political themes through natural disasters to entertainment and sports; from breaking news events through foreseeable, regularly occurring activities to channels for continuous thematic discussion; from local issues to global events; from events which unfolded over the span of a few hours to themes which were discussed for close to a year (and remain active beyond the timespan covered in our analysis); and from activities which involve only a relatively small subset of the global *Twitter* userbase, measuring in the thousands, to events which attracted the participation of close to two million unique users or generated more than six million tweets (see table 1).

Real-time data collection for these datasets generally commenced as the hashtags related to specific themes and events became prominent on *Twitter*; especially in the case of acute crisis events, this required researchers to react speedily as news of these crises (for example, natural disasters such as the Christchurch earthquake or the Japanese tsunami) broke, to rapidly determine the most prominent hashtags (*#eqnz*, *#tsunami*), and to add those hashtags to the existing *yourTwapperkeeper* installations for tracking. In other cases (such as *#royalwedding*, *#eurovision*, or *#ausvotes*), hashtags were foreseeable prior to the event, and could be added to the tracker in advance. In each case, however, we have further determined appropriate start and end points for the data timeframes to be considered in the present article, in order to focus on the key period of activity for each hashtag or keyword: for natural disaster events, this usually means limiting the analysis to the first days or weeks after the initial disaster event; for election-related discussion, the days around election day itself. The specific timeframes chosen for each hashtag or keyword are outlined in table 1. It is necessary to draw on this disparate collection of datasets for our analysis in order to detect any patterns of *Twitter* use which persist even in spite of such marked differences between individual cases.

In addition to the hashtag datasets (which contain only those tweets about a topic that were explicitly hashtagged – e.g. #tsunami for the March 2011 tsunami), we also include five *keyword* datasets: these were gathered by capturing all tweets that contained only the specific keyword (e.g. ‘tsunami’), regardless of whether or not the ‘#’ symbol was prefixed to that term. We discuss these datasets in more detail below; our aim in including them in the following analysis is to examine whether there are any indications that beyond the use of dedicated hashtags, topical communication patterns on *Twitter* may follow similar principles as we outline them for deliberately hashtagged exchanges.

User Activity Metrics

Given the divergence in the number of tweets and unique users for each dataset, it is first useful to compare the relative prominence of the leading user groups across these cases. Fig. 1 presents the relative amount of tweets contributed by each of the three user groups we have outlined above: lead users (top 1% most active users), highly active users (next 9% of active users), and least active users (the remaining 90% of users: the ‘long tail’ of the userbase).

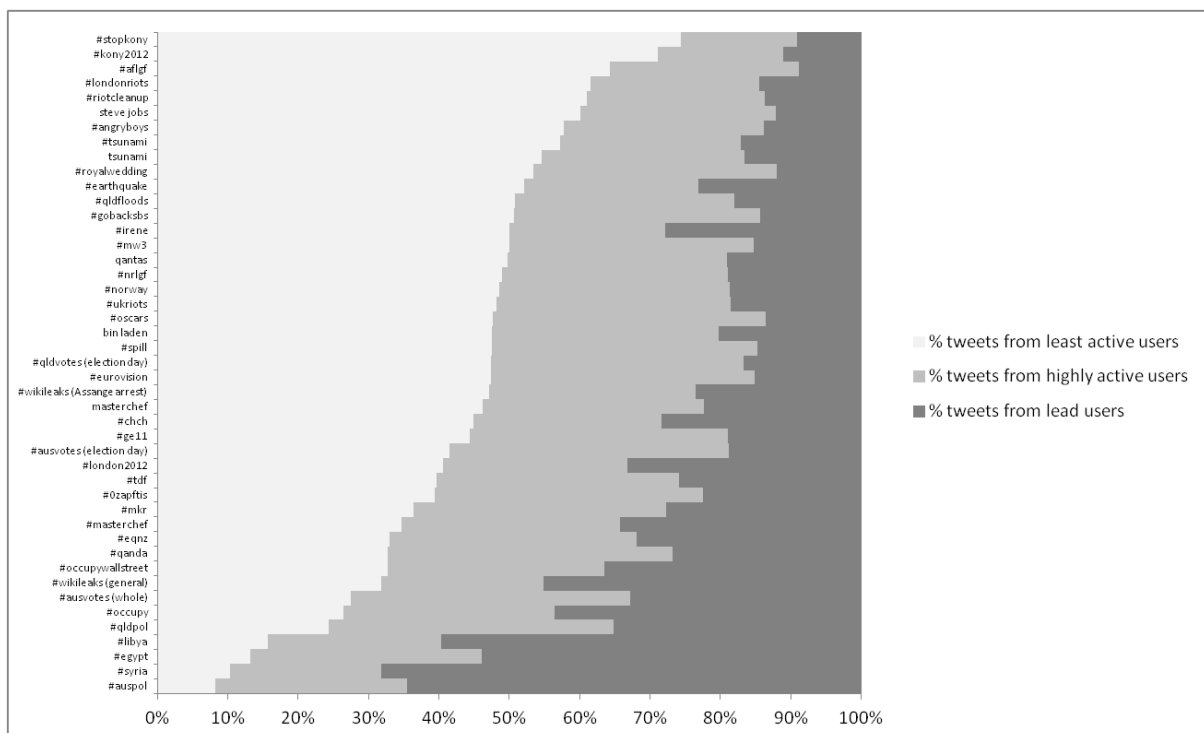


Fig. 1: Relative contributions from the three user groups

Clear distinctions between the cases examined here emerge from this analysis. Roughly half of the cases are comparatively dominated by the two most active user groups, who (in combination) contribute 50% or more of the total volume of tweets. For a smaller number of cases, that percentage grows to well above 70%: here, the 'long tail' of least active users remains largely silent, while any meaningful exchanges take place mainly within a dedicated in-group of highly active participants.

It is notable in this context that the hashtags which feature the most active groups of leading users are generally also those which cover the longest timeframes: in our comparison, the ten datasets which see the fewest tweets from the least active user group are #syria, #egypt, and #libya (each of which attracted hundreds of thousands of participants and was active throughout 2011); #auspol, a standing hashtag for the discussion of Australian federal politics with a small but highly active contributor community, and #ausvotes (for the discussion of the 2010 Australian federal election); #occupy, #occupywallstreet, and #wikileaks (which serve as key distribution tools for information about global political protest and counterculture movements, over the long term, and attract hundreds of thousands of participants); and #qanda (the hashtag promoted by the Australian Broadcasting Corporation for its weekly political talkshow *Q&A*).

By contrast, those hashtags which feature the most active 'long tails' of contributors in our analysis are also those which unfold over comparatively short timeframes: they include scheduled media events such as #royalwedding (the 29 April 2011 wedding between Prince William and Kate Middleton) or #aflgf (the 1 Oct. 2011 Australian Football League grand final); breaking news events such as the March 2011 Japanese tsunami (in both its hashtag and keyword variants), the August 2011 London riots and the subsequent #riotcleanup initiative organised by affected communities, and the death of Apple leader Steve Jobs (which we captured in a keyword dataset centred on mentions of 'Steve Jobs'); and short-lived viral marketing campaigns such as the initiative to bring Ugandan warlord Joseph Kony to justice (under the hashtags #kony2012 and #stopkony).

From these observations, we suggest that the relative prominence of leading user groups in a hashtag conversation is related to the overall longevity of the hashtag itself (the amount of time during which it was significantly active, and during which we gathered tweets for it): in a comparatively new hashtag, more striated community structures have not yet had a chance to crystallise, while in a long-lived hashtag it is perhaps logical that committed long-term contributors will emerge as lead users as more casual participants come and go. Fig. 2, which plots the longevity of hashtags against the relative contribution made by the 90% least active users, supports this finding. Although the possibility is intriguing, it should also be noted that our

evidence does not permit us to establish any causal relations between these factors, however: from our data alone it is impossible to determine whether hashtags persist for the longer term *because* a strong group of lead users keeps them going, or whether on the other hand these lead users inevitably emerge if a hashtag continues to remain active for a long enough time.




Fig. 2: Hashtag longevity compared to percentage of tweets contributed by 90% least active users

(size of data points indicates total number of tweets for each hashtag/keyword case)

Tweet Type Metrics

The hashtag datasets examined here also differ widely in their underlying communicative practices, as fig. 3 shows. Here, we examine the relative presence of the three key tweet types we have outlined above (original

tweets, genuine @replies, and retweets), as well as the occurrence of URLs in any such tweets. It is again obvious that there are distinct differences in communicative patterns between hashtags: most obviously, a small number of cases consist overwhelmingly (at 65% or above) of original tweets which neither mention nor reply to other users. These cases (and indeed, all hashtags with more than 50% original tweets) are also marked by the relative absence of URLs in tweets; the vast majority of this group of hashtags contain URLs in less than 20% of all tweets, while the average percentage of tweets with URLs for the remainder of our hashtags is close to 50%.




Fig. 3: Relative percentages of different tweets types across all hashtags

(note that the percentage of URLs is shown on a separate scale, as URLs can occur across all three tweet types)

It is notable in this context that hashtags which exhibit a large percentage of original tweets share a number of key contextual characteristics. For the most part, these hashtags relate to major media events, ranging from internationally televised entertainment broadcasts (#eurovision, #royalwedding, #oscar) through important sporting events (#tdf for the Tour de France, #aflgf for the Australian Football League Grand Final, #nrlgf for the Australian National Rugby League grand final) to popular daily or weekly television shows (the Australian reality TV programmes *Masterchef* and *My Kitchen Rules* – #mkr –, the sitcom *Angry Boys*, or the political talkshow *Q&A*). Other Australian political events – such as #spill for the 2010 partyroom challenge

against Prime Minister Kevin Rudd, or the #ausvotes discussion around election day 2010 – also fit this model, as they were (or for #spill, rapidly became) major media events in their own right.

On the other hand, hashtags which saw a substantial amount of retweeting, and comparatively few original tweets, largely fall into a category which may be best described as ‘breaking news’ or ‘rapid information dissemination’: they include, most obviously, many hashtags related to natural disasters from #eqnz (and the alternative #chch, short for Christchurch) through #earthquake and #tsunami (both relating specifically to the March 2011 event in Japan) to #qldfloods, as well as to civil unrest (from #libya through #occupywallstreet to the London #riotcleanup). Additional examples for this category are #stopkony (an orchestrated viral marketing campaign which to some extent behaved like a crisis event) and #0zapftis (a scandal around a Trojan horse virus developed by German intelligence services for covert investigation purposes). Generally, such retweet-heavy hashtags also contain a substantial number of tweets with URLs: on average, half of all tweets in hashtags with more than 50% retweets contain URLs.

Between these two key metrics, patterns in genuine @replies for each case move somewhat more randomly, and this category generally accounts only for a relatively small percentage of tweets in each dataset (with @replies constituting more than 41% of all tweets, #auspol is the one major exception to this rule; this may be related to the very well-established, dominant group of lead users in this case). In this context, a limitation of our hashtag-based *Twitter* research approach must be noted: as users respond to hashtagged tweets, they frequently do not again include the hashtag in their @replies, and such non-hashtagged replies are therefore not included in our datasets. Those users who *do* include a hashtag in their @replies, by contrast, usually do so explicitly to make their responses visible to a wider audience again; hashtagged @reply conversations are in essence performed in front of a larger public, in other words, but constitute a special case. For this reason, the following discussion largely focusses on the complementary metrics of original tweets and retweets only, as well as on the presence of URLs.

Towards a Typology of Hashtags

These observations enable the development of a tentative typology of hashtag events, based on the activity metrics which are observable in each case. For the datasets we have examined here, fig. 4 plots the

percentage of URLs in tweets against the percentage of retweets in the overall dataset, and indicates the combined contribution of lead and highly active users through the size of each data point:




Fig. 4: percentage of URLs in tweets vs. percentage of retweets amongst all tweets

(size of data points shows combined contribution of lead users and highly active users)

On this graph, two distinct clusters of hashtag cases emerge. At the centre of fig. 4 is a cluster which mainly contains hashtags relating to crises and other breaking news events; these range from natural disasters to political protests and civil unrest. This group of hashtags is characterised by both substantial posting of links to further information, and significant retweeting activity; we suggest, therefore, that it is largely centred around a shared practice of *gategatching* (Bruns, 2005): the collaborative identification, sharing, and passing-along of situationally relevant information, here especially in the context of “acute events” (Burgess & Crawford, 2011).

Indeed, it is notable for this cluster of hashtags that most natural disasters – except for the #irene hashtag, for the 2011 Hurricane Irene – are positioned towards the top of the cluster (indicating an especially high percentage of retweets); this may indicate a widespread desire of users to help in sharing key emergency information, and a limited interest in posting comments or other statements about the unfolding event, while political crises attract a comparatively higher number of such comments in the form of original tweets and @replies. By contrast, on the far right of the cluster we find a number of hashtags that are related to political protest movements (#wikileaks, #occupy, and #occupywallstreet)²: their positioning indicates a substantial percentage of URLs being shared through the hashtag, pointing perhaps to contributors' perception of these themes as countercultural issues which are under- or misreported by mainstream media and require constant support through online social networks. (In this context, it is notable that #occupywallstreet – which deals more centrally with the struggle between protesters and law enforcement in New York – and the subset of #wikileaks activity around the arrest of Julian Assange share more similarities with the overall crisis cluster than the longer-term #wikileaks and #occupy hashtags.)

In addition to these hashtags, we have also included metrics for keyword datasets covering mentions of Steve Jobs following his death, for Osama bin Laden after his death in a raid on his Abbottabad compound, and for the Australian airline Qantas during a global grounding of all flights by management in response to an industrial dispute. Of these, the activity metrics for Jobs and bin Laden show strong similarities to other crisis events, pointing to similar gatewatching and news-sharing activities in the context of these breaking news events. The Qantas event behaves somewhat differently, due to a comparatively lower percentage of retweets (and thus a larger number of tweets making original comments) – this is in keeping with the significant political implications of the event, beyond the international air transport crisis it caused. Further, we also include the keyword dataset for 'tsunami' in addition to the hashtag #tsunami, and find a comparatively smaller percentage of retweets for the keyword case; this indicates, not unexpectedly, that hashtagged tweets are more likely to be found and retweeted than non-hashtagged messages, but also points to the likelihood that overall *Twitter* activity patterns around crisis events, beyond their hashtagged core, are broadly similar to those for the crisis hashtag itself.

² We consider such protest movements, which largely remain focussed on leading nations in the west, to be distinct from the civil unrest in Libya, Egypt, or Syria, whose hashtags are located at the centre of the cluster.

Finally, we find the #kony2012 hashtag at the centre of the crisis cluster, while its cousin #stopkony is present as an outlier in the overall graph, with a very substantial percentage of retweets but comparatively few URLs. Further analysis must determine the reasons for the latter hashtag's divergent behaviour, but it appears sensible that #kony2012 – a campaign designed to be disseminated virally and with reference to further information on the campaign Website, and videos on *YouTube* – would show similar tendencies to the crisis hashtags; in essence, we may understand #kony2012 (and similar viral campaigns) as a deliberately 'manufactured' crisis.

A second distinct cluster of hashtags is located in the bottom left quadrant of the graph: the hashtags assembled here are characterised by a very low percentage of URLs in each dataset, as well as a comparatively low percentage of retweets; put differently, these hashtags are mainly used to post original tweets and a limited amount of @replies, with few references to additional information outside of *Twitter*. Where the hashtags assembled in the acute events cluster are largely concerned with information sharing, therefore, the hashtags in this second cluster are focussed on original commentary.

Further, the majority of these hashtags are clearly related to mainstream media events, as we have already seen in the discussion of fig. 3; we therefore interpret these hashtags as cases in which *Twitter* functions as a backchannel for live events (especially as they are broadcast by national and international television). These hashtags, in other words, support a shared experience of 'audiencing' (e.g. Fiske, 1992): of talking back at the television (or the live event), along with thousands of other viewers. This sense of temporary, imagined community persists even if – as our data show – actual direct interaction between users through hashtagged @replies and retweets remains relatively rare; it may be sufficient to observe the stream of hashtagged comments, even without engaging with and replying to them. (Such a sense of community is further enhanced, of course, if – as is increasingly common practice – television shows include selected tweets from the hashtag stream in an on-screen ticker.)

In this cluster, too, further subdivisions can be observed: interaction through retweets is lowest for sporting and other entertainment events, while political themes attract a somewhat larger percentage of retweets – the #qldvotes, #ausvotes, and #ge11 (for the 2011 Irish general election) hashtags on their respective election days, as well as #spill (for the 2010 Australian political leadership crisis), are located towards the top of this cluster. *Go Back to Where You Came From* (#gobacksbs), while in principle a reality TV show, must similarly be included here, as it thematised the highly controversial theme of asylum seeker policy

in Australia; by contrast, it is notable that the hashtag for the overtly political talkshow *Q&A* does not show a retweet rate which is comparable to other political backchannel cases, for reasons which remain as yet unclear. Finally, while not immediately connected to any one mainstream media channel or show, the well-established #auspol hashtag, hosting a continuous discussion of Australian political events, appears to operate much like the other hashtags within this cluster; it may be understood, therefore, as an aggregate backchannel to mainstream political news reporting in the country, rather than as collective effort to engage in gatewatching or other citizen journalism activities.

In discussing both these clusters, it is important to note that they emerge from our analysis even in spite of the widely divergent timeframes for these individual hashtag cases (ranging from hours and days to close to a year), and the varying sizes of the hashtags' userbases (from less than 10,000 to more than 2 million participants). This points to the fact that these patterns of activity reflect standard uses of *Twitter*, which participants engage in as the theme and purpose of the hashtag demands it; it appears that the backchannel to a minor television series does not operate much differently from that for a global media event, and the response to a natural disaster does not change substantially as a greater number of people are affected.

Similarly, while (as we have seen above) the dominance of leading user groups appears to be related to the longevity of a hashtag, the activity patterns which we have observed here do not depend on the activities of that leadership group alone, as fig. 5 demonstrates. It explores the presence of any correlations between the combined contributions made by the top 10% of most active users, and the percentage of URLs in the total dataset, and finds no significant connections between these metrics; a corresponding graph comparing the contributions of leading users and the percentage of retweets would similarly yield no correlations.




Fig. 5: percentage of tweets contributed by lead and highly active users vs. percentage of URLs in tweets

(size of data points shows total number of tweets per hashtag)

This is an important observation, as it shows activity patterns in a hashtag (as measured by the percentage of retweets or URLs) to be independent of the internal make-up of the hashtag community (as measured through the 1/9/90 distinction between user groups). In fig. 5, the group of backchannel hashtags at the bottom of the graph remains clearly separate from the group of acute event hashtags at the centre; the same is true for the percentage of retweets in each hashtag – leading and peripheral users may be different in many respects, but their understanding of acute events and shared audiencing experiences appears to be similar nonetheless.

Conclusion

What emerges from this wide-ranging comparison of participation patterns across a diverse collection of hashtag datasets is that *Twitter* activities, especially around defined themes and events, are far from random, but instead appear to be governed by a number of standard practices. Of these, the practices of gatewatching and audiencing are most obviously visible in our analysis, and relate clearly to the underlying themes of the hashtags we have examined: a standard response to the emergence of breaking news and other acute events is the tendency to find, share, and reshare relevant information, resulting in a high rate of URLs and retweets, while for live, mainstream media events *Twitter* acts as a backchannel, containing mainly original commentary which does not engage with the tweets of others or provide a substantial number of links to further information.

For any research dealing with *Twitter* data, it must be noted that due to the vagaries of working with the *Twitter* API itself, as well as because of unavoidable disruptions caused by regular maintenance to the university servers on which *yourTwapperkeeper* was run, the datasets thus created does not constitute an entirely comprehensive corpus of *all* tweets that included the specific keywords; indeed, it is true that unless the *Twitter* API can be trusted to deliver all matching tweets without disruption, no study of *Twitter* which uses these processes can possibly achieve 100% accuracy. Further, as the API is the only access point to large-scale *Twitter* data which is available to researchers outside of *Twitter* itself, there is no opportunity to independently verify the quality of the dataset. This is a necessary and unavoidable limitation which does not invalidate the findings of studies such as ours, however; any sufficiently complex system of communication will suffer from a certain level of message loss.

Furthermore, it has to be considered that the types of topical hashtags addressed here are not the only ones which may be observed on *Twitter*; further research is required to establish similar metrics for a wider range of *Twitter* events and to compare them with the metrics we have presented here. For example, it may well be possible that a greater number of counterculture and protest politics hashtags may exhibit similar patterns to what we have already observed for #wikileaks and #occupy, forming their own distinct cluster of cases; the intra-cluster distinctions we have noted for both the acute events and the backchannel cluster may also turn out to be more pronounced as more examples are added. Further, it must also be remembered that the uses of *Twitter* continue to evolve, especially also as a consequence of each major new event – while in

combination, our datasets cover a period of some two years, it remains to be seen whether future events will continue to show similar patterns of activity.

At the same time, if these patterns are indeed consistent across a larger number of cases and for the longer term, then our findings may also open up possibilities to operationalise them in the detection of new *Twitter* events. It may be possible, for example, to distinguish new acute events from other hashtags by calculating their activity metrics; this could be of use for media monitoring and emergency operations as it would point to the emergence of crisis events purely on the basis of activity metrics, even if relevant keywords have not yet been identified. Further, if – as our examination of a handful of keyword archives appears to suggest – non-hashtagged keywords behave largely similarly to their related hashtags, this may support the identification of acute events (and their distinction from other trending topics) before users have even agreed on a standard hashtag to adopt.

Finally, our research points to the potential of understanding patterns of *Twitter* activity at large scale, beyond (but building on) the study of individual communicative events. Such ‘big data’ research (boyd & Crawford 2011), drawing on comprehensive access to user activity data through platform APIs, remains in its infancy but is set to generate significant new opportunities for researchers in the humanities and allied disciplines. Current work on *Twitter*, such as the research presented here, will be able to be usefully combined and compared with studies of other (social) media platforms in order to develop a more comprehensive and detailed picture of information and communication flows in society, in turn providing the basis both for a more sophisticated understanding of the place of social media in society, and of potential points of leverage for relevant institutions (for example governments, media, or emergency services) as they seek to engage with and influence such information flows.

But while this article has focussed almost exclusively on the examination of large-scale, quantitative patterns in *Twitter* datasets of considerable size, this should not be misunderstood to privilege such quantitative research over other approaches. Rather, we close by noting the substantial opportunities for qualitative and combined quantitative/qualitative research which also exist in this field (e.g. Krüger *et al.*, 2012). To begin with, the quantitative approaches to understanding communicative patterns on *Twitter* which we have introduced here (for a more detailed discussion, also see Bruns & Stieglitz, forthcoming) provide an opportunity to pinpoint specific areas for further, detailed, qualitative investigation. A focus on the communicative activities of representatives of the different groups of lead, highly active, and least active users

which we have introduced above enables an examination of a variety of distinct tweeting styles within the same hashtag exchange, for example; representatives of each group could also be studied in much greater detail through in-depth ethnographic work, or engaged through survey or interview techniques to better understand these diverse approaches to using *Twitter* in specific communicative contexts.

What must be noted in this context is that the process of generating overall headline metrics for each of these datasets does not destroy the datasets themselves, which remain available for much closer, tweet-by-tweet analysis. So, for example, for datasets which follow the overall gatewatching pattern of collaborative sourcing and sharing information which appears to be common to crisis events, a possible avenue for further research is the qualitative (or mixed-methods) study of how these patterns emerge in each case, and whether these processes of emergence generally follow similar steps. Potential questions to be addressed here include how groups of lead users crystallise from the early participant base; how they come to structure their activities as the acute event unfolds; and how common principles and shared understandings of how to engage with the event are established in each case. Such questions (and similar questions which apply to the non-crisis events amongst the datasets we have examined here) may be addressed, *inter alia*, through a close, qualitative reading of the relevant tweets in the dataset, through ethnographic studies of user communities, or through other methods drawn from media, cultural and communication studies, anthropology, or the social sciences. The metrics which we have outlined here, and their utilisation for a birds-eye comparison of large-scale communicative events on *Twitter*, are intended to serve as a useful, necessary starting-point for such further research endeavours.

References

- boyd, d., & Crawford, K. (2011). Six provocations for big data. Paper presented at A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society, Oxford Internet Institute, 21 Sep. 2011.
- , Golder, S., & Lotan, G. (2010, 6 Jan.). Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. HICSS-43, Kauai, Hawai'i. <http://www.danah.org/papers/TweetTweetRetweet.pdf> (accessed 13 May 2012).
- Bruns, A. (2005). *Gatewatching: Collaborative Online News Production*. New York: Peter Lang.
- . (2012). Ad Hoc innovation by users of social networks: The case of Twitter. *ZSI Discussion Paper 16* (2012). <https://www.zsi.at/object/publication/2186> (accessed 13 May 2012).
- , & Burgess, J. (2011a). #ausvotes: How Twitter covered the 2010 Australian federal election. *Communication, Politics & Culture*, 44(2), 37-56.

- , & Burgess, J. (2011b). The use of Twitter Hashtags in the Formation of *Ad Hoc* Publics. Paper presented at the 6th European Consortium for Political Research General Conference, University of Iceland, Reykjavík, 25-27 Aug. 2011. <http://eprints.qut.edu.au/46515/> (accessed 3 Feb. 2012).
- , & Liang, E. (2012). Tools and methods for capturing Twitter data during natural disasters. *First Monday*, 17(4). <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3937/3193> (accessed 13 May 2012).
- , & Stieglitz, S. (forthcoming 2013). Metrics for Understanding Communication on Twitter. In K. Weller et al., eds., *Twitter and Society*. New York: Peter Lang.
- , Burgess, J., Crawford, K., & Shaw, F. (2012). *#qldfloods and @QPSMedia: Crisis communication on Twitter in the 2011 South East Queensland floods*. Brisbane: ARC Centre of Excellence for Creative Industries and Innovation. <http://cci.edu.au/floodsreport.pdf> (accessed 13 May 2012).
- , Highfield, T., & Burgess, J. (forthcoming). The Arab Spring and its social media audiences: English and Arabic Twitter users and their networks. *American Behavioural Scientist*.
- Burgess, J., & Crawford, K. (2011). Social media and the theory of the acute event. Paper presented at Internet Research 12.0 – Performance and Participation, Seattle, October 2011.
- Cheng, J., Sun, A., Hu, D., & Zeng, D. (2011). An information diffusion-based recommendation framework for micro blogging. *Journal of the Association of Information Systems* 12(7), 463-486.
- Cheong, F. & Cheong, C. (2011). Social media data mining: A social network analysis of tweets during the 2010-2011 Australian floods. *Proceedings of the Pacific Asia Conference on Information Systems*, Paper 46.
- Christensen, C. (2011). Twitter revolutions? Addressing social media and dissent. *Communication Review* 14(3), 155-157.
- Conover, M. D., Ratkiewicz, J., Francisco, M., Gonalves, B., Flammini, A., & Menczer, F. (2011). Political polarization on Twitter. *Proceedings of the 5th International Conference on Weblogs and Social Media*.
- Deller, R. (2011). Twittering on: Audience research and participation using Twitter. *Participations* 8(1). <http://www.participations.org/Volume%208/Issue%201/deller.htm> (accessed 13 May 2012).
- Dröge, E., Maghferat, P., Puschmann, C., Verbina, J., & Weller, K. (2011). Konferenz-Tweets. Ein Ansatz zur Analyse der Twitter-Kommunikation bei wissenschaftlichen Konferenzen. *Proceedings of the 12th International Symposium for Information Science*. Verlag Werner Hülsbusch. <http://ynada.com/pubs/isi2010.pdf> (accessed 13 May 2012).
- Fiske, J. (1992). Audiencing: A cultural studies approach to watching television. *Poetics*, 21(4), 345-359.
- Hughes, A., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management* 6(3-4), 248-60.
- Krishnamurthy, B., Phillipa, G., & Arlitt, M. (2008). A few chirps about Twitter. *Proceedings of the first workshop on online social networks (WOSN '08)*.
- Krüger, N., Stieglitz, S. & Potthoff, T., (2012). Brand Communication on Twitter – A Case Study on Adidas. *Proceedings of the 16th Pacific Asia Conference on Information Systems*, Paper 343.
- Larsson, A.O., & Moe, H. (2011). Studying political microblogging: Twitter users in the 2010 Swedish election campaign. *New Media & Society*. doi:10.1177/1461444811422894
- Lotan, G., Ananny, M., Gaffney, D., & boyd, d. (2011). The revolutions were tweeted: Information flows during the 2011 Tunisian and Egyptian revolutions. *International Journal of Communication* 5, 1375-1405.
- Marwick, A.E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1), 114-133.
- Mendoza, M., Poblete, B., & Castillo, C. (2010). Twitter under crisis: Can we trust what we RT? *1st Workshop on Social Media Analytics (SOMA '10)*, Washington, DC.
- Palen, L., Starbird, K., Vieweg, S., & Hughes, A. (2010). Twitter-based information distribution during the 2009 Red River Valley flood threat. *Bulletin of the American Society for Information Science and Technology*, 36(5), 13-17.
- Papacharissi, Z. (2011). Conclusion: A networked self. In Z. Papacharissi (ed.), *A Networked Self: Identity, Community and Culture on Social Network Sites*. New York: Routledge, 304-318.

- Park, J., Cha, M., Kim, H., & Jeong, J. (2011). Managing bad news in social media: A case study on Domino's Pizza crisis. *Proceedings of the ICWSM 2012*.
- Suh, B., Hong, L., Pirolli, P., & Chi, E. (2010). Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network. *Proceedings of the IEEE International Conference on Social Computing*.
- Stieglitz, S., & Dang-Xuan, L. (2012). Political communication and influence through microblogging – An empirical analysis of sentiment in Twitter messages and retweet behavior. *Proceedings of the 45th Hawaii International Conference on System Sciences*, 3500-3509.
- Stieglitz, S., & Dang-Xuan, L. (forthcoming). Social Media and Political Communication – A Social Media Analytics Framework. *Social Network Analysis and Mining*, Springer.
- Stieglitz, S., & Krüger, N. (2011): Analysis of sentiments in corporate Twitter communication – A case study on an issue of Toyota. *Proceedings of the 22nd Australasian Conference on Information Systems*, Paper 103.
- Tedjamulia, S. J. J., Dean, D. L., Olsen, D. R., & Albrecht, C. C. (2005). Motivating content contributions to online communities: Toward a more comprehensive theory. *Proceedings of the 38th Annual Hawaii International Conference on System Science*.
- Tumasjan, A., Sprenger, T., Sandner, P., & Welp, L. (2011). Election forecasts with Twitter: How 140 characters reflect the political landscape. *Social Science Computer Review* 29(4), 402-418.
- Twitter (2012, 21 Mar.). Twitter turns six. <http://blog.twitter.com/2012/03/twitter-turns-six.html> (accessed 13 May 2012).
- Vis, F. (2012, 24 Jan.). Reading the riots on Twitter: Who tweeted the riots? *Researching Social Media*. <http://researchingsocialmedia.org/2012/01/24/reading-the-riots-on-twitter-who-tweeted-the-riots/> (accessed 13 May. 2012).
- Weller, K., Dröge, E., & Puschmann, C. (2011). Citation analysis in Twitter: Approaches for defining and measuring information flows within tweets during scientific conferences. #MSM2011 – 1st Workshop on Making Sense of Microposts. <http://files.ynada.com/papers/msm2011.pdf> (accessed 13 May 2012).
- Yardi, S., & boyd, d. (2010). Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of Science, Technology and Society* 20, 1-8.
- yourTwapperkeeper*. (2012). <https://github.com/jobrieniii/yourTwapperKeeper> (accessed 25 Aug. 2011).

Hashtag / Keyword	Description	Theme	Timeframe	Notes on timeframe	Total unique users	Total tweets
#0zapftis	Scandal around trojan horse virus developed by German intelligence service	Politics	11-31 Oct. 2011	Three weeks following scandal	6716	26158
#aflgf	Australian Football League 2011 grand final	Sports	1-2 Oct. 2011	Matchday and following day	3793	6135
#angryboys	<i>Angry Boys</i> : popular weekly TV sitcom on Australian free-to-air television	Television	12 May to 31 July 2011	Full season	30121	63333
#auspol	Australian politics (general discussion)	Politics	8 Feb. to 8 Dec. 2011	Eight months	26290	854019
#ausvotes	Australian federal election 2010	Politics	20-22 Aug. 2010	Three days around election day	36286	415511
#chch	February 2011 earthquake in Christchurch, New Zealand	Natural disaster	22-28 Feb. 2011	First week after earthquake (initial alternative to #eqnz)	9688	24400
#earthquake	March 2011 earthquake and tsunami in Japan	Natural disaster	11-24 Mar. 2011	First two weeks after earthquake	183794	358737
#egypt	Arab Spring protests	Political unrest	26 Feb. to 26 Nov. 2011	Nine month period since first major protests	281978	6277782
#eqnz	February 2011 earthquake in Christchurch, New Zealand	Natural disaster	22 Feb. to 7 Mar. 2011	First two weeks after earthquake	37635	156940
#eurovision	Eurovision Song Contest 2011	Television	9-15 May 2011	Semi-finals and and finals broadcasts on 10/12/14 May	137745	520543
#ge11	Irish general election	Politics	26 Feb. 2011	Election day	6151	28468
#gobacksbs	<i>Go Back to Where You Came From</i> : weekly reality TV show with political connotations on Australian free-to-air television	Television	21 June to 4 July 2011	Second half of season	8691	29009
#irene	Hurricane Irene along the East Coast of the U.S.	Natural disaster	27 Aug. to 17 Sep. 2011	Three weeks following first impact	37891	64315
#kony2012	Viral campaign to arrest warlord Joseph Kony	Politics	8-21 Mar. 2012	First two weeks of campaign	80874	101425
#libya	Arab Spring protests	Political unrest	26 Feb. to 26 Nov. 2011	Nine month period since first major protests	363489	3825272
#londonriots	Violent riots in London and the UK	Crisis	8-21 Aug. 2011	First two weeks after riots	127631	212213
#masterchef	<i>Masterchef</i> : popular weekly reality TV show on Australian free-to-air television	Television	1 May to 8 Aug. 2011	Whole season	54117	210773
#mkr	<i>My Kitchen Rules</i> : popular weekly reality TV show on Australian free-to-air television	Television	13 Feb. to 31 Mar. 2012	Final 27 episodes	12671	63866
#mw3	<i>Modern Warfare 3</i> : popular computer game	Net culture	1-30 Nov. 2011	One month around official launch	207858	413922
#norway	Right-wing terrorist attacks in Oslo and Utøya	Crisis	24 July to 9 Aug. 2011	First two weeks after attacks	38224	63244
#nrlgf	Australian National Rugby League 2011 grand final	Sports	1-2 Oct. 2011	Build-up and matchday	2049	4182
#occupy	Global Occupy protests	Political protests	19 Dec. 2011 to 19 Apr. 2012	Four months of protests	121952	560560
#occupywallstreet	Occupy protests in New York	Political protests	27 Sep. to 27 Nov. 2011	Three months at height of protests	234514	885174
#oscars	Academy Awards 2011	Entertainment	27 Feb. 2011	Event day	236103	639251
#qanda	<i>Q&A</i> : popular weekly political talkshow on Australian free-to-air television	Politics	21 Feb. to 21 Nov 2011	Whole season	246231	47131
#qldfloods	Major flooding in south-east Queensland	Natural disaster	10-16 Jan. 2011	First week of floods	15553	35658
#qldvotes	2012 Queensland state election	Politics	23-25 Mar. 2012	Three days around election day	5788	17456
#riotcleanup	Clean-up after violent riots in London and the UK	Crisis	8-21 Aug. 2011	First two weeks after riots	38511	53381
#royalwedding	Wedding between Prince William and Kate Middleton	Entertainment	29 Apr. 2011	Wedding day	492566	926527
#spill	Party room revolt against Australian Prime Minister Kevin Rudd	Politics	23-24 June 2010	First rumours and confirmation of party room vote	11309	46937
#stopkony	Viral campaign to arrest warlord Joseph Kony	Politics	8-21 Mar. 2012	First two weeks of campaign	117050	140958
#syria	Arab Spring protests	Political unrest	26 Mar. to 26 Nov. 2011	Eight months since first major protests	229030	5230025
#tdf	Tour de France 2011	Sports	4 July to 26 July 2011	Whole tour (except first two days)	94830	427467
#tsunami	March 2011 earthquake and tsunami in Japan	Natural disaster	11 Mar. to 11 Apr. 2011	First month after earthquake	529913	948640
#ukriots	Violent riots in London and the UK	Crisis	8-21 Aug. 2011	First two weeks after riots	61766	126664
#wikileaks	Political controversy	Politics	26 Feb. to 26 Nov. 2011 + 1-7 Sep. 2011	Nine months + period around Julian Assange arrest in the UK	119853 16930	422635 35451
bin laden	Death of Osama bin Laden	Politics	2 May to 2 June 2011	First month following bin Laden killing in Abbottabad	1868127	3987919
masterchef	<i>Masterchef</i> : popular weekly reality TV show on Australian free-to-air television	Television	1 May to 8 Aug. 2011	Whole season	238689	609714
qantas	Global grounding of flights by Qantas management in response to industrial action	Brand crisis	26 Oct. to 8 Nov. 2011	Two weeks around grounding crisis	42144	98636
steve jobs	Death of Apple founder Steve Jobs	Net culture	7 Oct. to 7 Nov. 2011	One month after death	403321	562411
tsunami	March 2011 earthquake and tsunami in Japan	Natural disaster	11 Mar. to 11 Apr. 2011	First month after earthquake	1936553	4246019

Table 1: overview of hashtag and keyword datasets used for the comparative analysis