Notes towards the Scientific Study of Public Communication on Twitter

Axel Bruns & Jean Burgess

ARC Centre of Excellence for Creative Industries and Innovation
Queensland University of Technology
Brisbane, Australia
a.bruns / je.burgess @ qut.edu.au
http://mappingonlinepublics.net/

Twitter *is now well-established as an important platform for real-time public communication.* Twitter *research continues to lag behind these developments, with many studies remaining focussed on individual case studies and utilising homegrown, ideosyncratic, non-repeatable and non-verifiable research methodologies. While the development of a full-blown 'science of* Twitter*' may remain illusory, it is nonetheless necessary to move beyond such individual scholarship and towards the development of more comprehensive, transferable, and rigorous tools and methods for the study of* Twitter*, at large scale and in close to real time.*

Introduction

Social media platforms like Twitter are playing a significant role in public communication – first among private individuals, and now increasingly also among media organisations, journalists, governments, and politicians in conversation or debate with their citizens, consumers and users. Researchers working with Twitter data at various levels of scale and complexity have already generated rich insights into the use of this social media platform: for personal communication, politics, journalism, crisis communication, and so on almost *ad infinitum* (see, for example, boyd et al, 2010; Crawford, 2009; Marwick & Boyd, 2011; Hermida, 2010). Within humanities and social science approaches to media and communication, as has been the case with the study of earlier new media technologies and forms (especially television and its audiences), we see an eclectic mix of methods, and radically different scales of analysis. In this paper, we argue for the importance of transferable methods, hence enabling meaningful comparative work across research teams and national traditions, if not disciplines; and hopefully for the more systematic coordination of multi-method approaches.

Our ability to compare the findings of Twitter research across individual case studies, in fact, is hindered by the lack of a standard set of communicative measures and metrics which may be applied in the analysis of Twitter datasets – if we are to pursue more 'scientific' approaches to Twitter research grounded in humanities and social science approaches to questions of media and communication, the development of such metrics will be an important contribution. In the following sections of the paper, we provide examples from our own research of how relatively simple metrics, particularly when used comparatively, at scale and over time, can yield analytically productive insights into longstanding questions of media and communication studies: Who are the main actors engaged around a topic or event? How might we think about the communicative and/or power relations among those actors? What are the main themes or frames associated with the social media communication around a topic or event?

This work is set against the much broader backdrop of what David Berry and others call the 'computational turn' – a 'third wave' of digital humanities which sees the shift from computational *tools* to a new computational *paradigm*, changing the ontologies and epistemologies of humanities research (Berry, 2012). Such a shift is represented, for example, by the work of Franco Moretti on large-scale, corpus-based literary analysis in the mid-2000s; as well as Richard Rogers' (2009) call to employ 'natively' digital methods to diagnose patterns of social change via the digital traces that can be gleaned via the Internet; rather than using the internet to carry out traditional social science or humanities enquiries – for Rogers, this is the distinction between 'virtual' and (natively) 'digital' methods. In what follows, we present the techniques resulting from our development of 'natively' digital approaches to communication via the use of the Twitter API, and discuss their applications.

Twitter Communication Metrics

The development of metrics for understanding public communication on Twitter naturally begins with a review of the datapoints the Twitter Application Programming Interface (API) already offers, directly or indirectly. In addition to the tweet text itself, the data and metadata which the API offers to describe a single tweet include a number of other key points of interest:

- **text:** contents of the tweet itself, in 140 characters or less
- **to_user_id:** numerical ID of the tweet recipient (for @replies)
- **from_user:** screen name of the tweet sender
- **id:** numerical ID of the tweet itself
- **from_user_id:** numerical ID of the tweet sender
- **iso_language_code:** code (e.g. en, de, fr, ...) of the sender's default language
- **source:** client software used to tweet (e.g. Web, Tweetdeck, ...)
- **profile_image_url:** URL of the tweet sender's profile picture
- **geo_type:** format of the sender's geographical coordinates
- **geo_coordinates_0:** first element of the geographical coordinates
- **geo_coordinates_1:** second element of the geographical coordinates
- **created_at:** tweet timestamp in human-readable format
- **time:** tweet timestamp as a numerical Unix timestamp

Further information can be extracted from the tweets themselves. An examination of the syntax of each tweet, for example, can reveal whether it should be classed as belonging to one of the following categories of communicative activity:

- **original tweets:** tweets which are neither @reply nor retweet
- **retweets:** tweets which contain RT *@user*… (or similar)
    - **unedited retweets:** retweets which start with RT *@user*…
    - **edited retweets:** retweets do not start with RT *@user*…
- **genuine @replies:** tweets which contain *@user*, but are not retweets
- **URL sharing:** tweets which contain URLs

(Any one tweet will be *either* an original tweet, retweet, *or* @reply, but tweets from each of these categories may *also* contain URLs.)

Although basic, such simple approaches to categorising tweets are already able to generate significant insights into the interaction patterns which may be observed for public communication on Twitter: our minute-by-minute examination of the #royalwedding hashtag which accompanied the 29 April 2011 British royal wedding, for example, clearly points to key moments in the day, such as the newlyweds' first public kiss on the balcony of Buckingham Palace, which resulted in a sharp spike in original tweets – expressing viewers' personal reactions to the moment – and simultaneous drops in retweeting, @replying, and link-sharing activities. Further, overall tweeting volumes also indicate the times at which major international television networks began and ended their coverage (see Bruns, 2011).

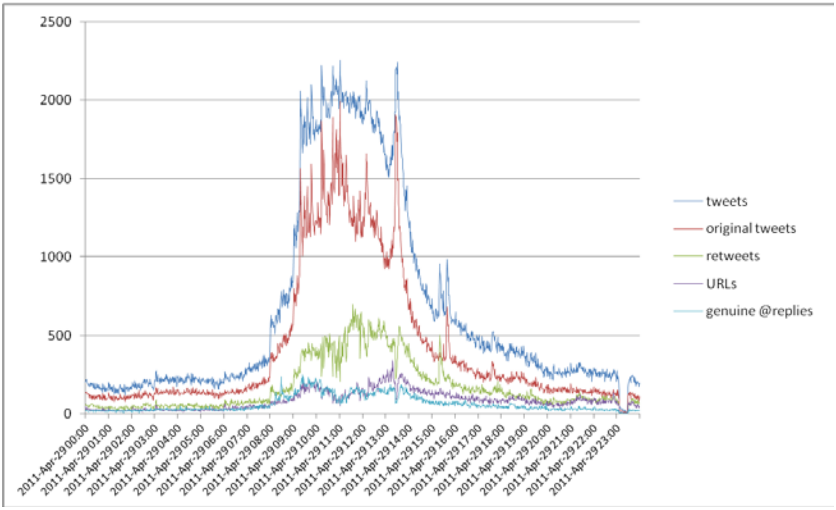

Fig. 1: minute-by-minute activity in the #royalwedding hashtag, 29 Apr. 2011 (times in GMT)

Once they are based on such standard metrics, such analyses of individual hashtagged events may then also be usefully compared across a range of different events, to identify shared or divergent patterns between activities of the same time. Bruns & Stieglitz (2012, under review) do so for a large number of hashtag datasets, and detect clear correlations between the wider communicative context within which specific hashtags operate, and the communicative patterns which may be observed within these hashtags themselves:

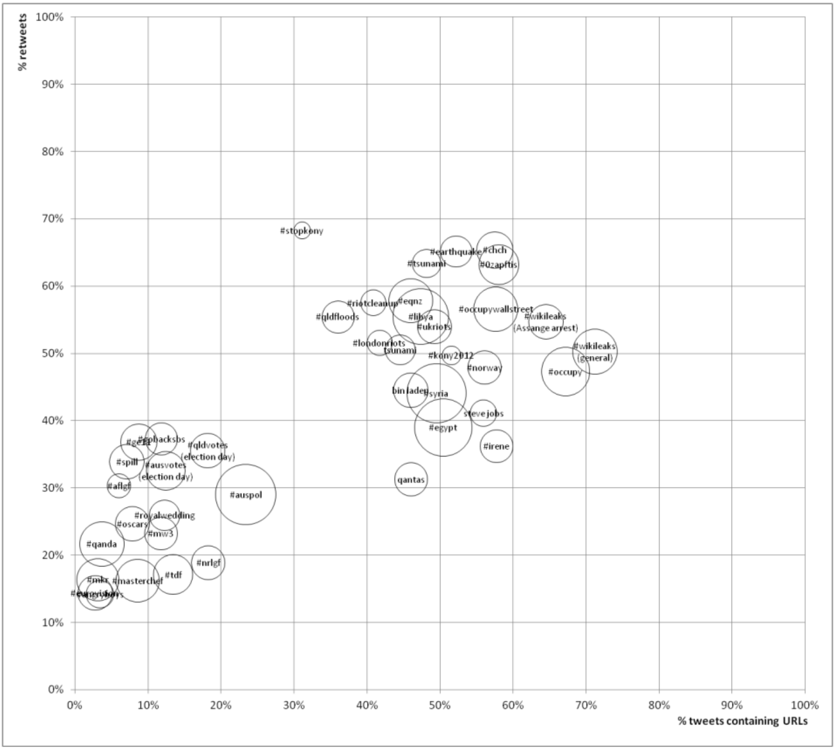

Fig. 2: percentage of URLs in tweets vs. percentage of retweets amongst all tweets (size of data points shows level of activity by leading 10% of users)

Their analysis of some 40 different hashtag events points to two clear, divergent patterns of hashtag activity: one the one hand, hashtags which are associated with breaking, unforeseen news events and crises (#egypt, #londonriots, #qldfloods) are characterised by a substantial level both of tweets containing URLs, and of retweets; user practices here can be described as a form of gatewatching (Bruns, 2005), with users actively seeking out and sharing information about the event at hand as it unfolds. On the other hand, a second cluster of hashtags contain very few URLs and a similarly smaller number of retweets: these hashtags are largely associated with widely televised, foreseeable events from sports through popular culture to election nights (#tdf, #oscars, #ausvotes), and users are mainly contributing by sending original tweets and engaging with one another through @replies.

To date, this analysis covers only a relatively small number of hashtags which relate to major events; it is entirely possible, therefore, that the addition of further metrics for a broader range of hashtags – denoting long-term communities of interest (e.g. #phdchat), more generic themes (#socmed), or even emotional responses (#headdesk) – might lead to the identification of additional types of hashtag use. Such work is only possible, of course, if standard metrics are applied to the study of such further communicative events on Twitter.

Twitter User Metrics

In addition to the development of such metrics for the description of communicative patterns in hashtagged conversations, additional standardised measures may also be established to examine the make-up and activities of the user communities – the *ad hoc* publics (Bruns & Burgess, 2011) – which form around such hashtags. In the first place, it is possible to use the distinction of tweet types which we have outlined above to describe the tweeting profile of each participating user: to examine, for example, the balance between original tweets, @replies, and retweets they have sent, and to correlate this with the number of @replies and retweets they have received in turn. Such analysis may be used, for instance, to distinguish accounts which merely retweet other users' messages, or post their own, from those which genuinely engage with others by @replying.

At hashtag level, however, such metrics may also be examined in connection with other communicative patterns. Central to such analysis is a further distinction of the hashtag community into its more or less active components: based, for example, on a simple division of the total contributor base for a hashtag into its leading, most active tweeters and other, less active groups, it becomes possible to determine the extent to which a small number of highly active participants dominate exchanges, and to examine differences in tweeting patterns across these groups of more or less active users. Our analysis of the well-established #auspol hashtag for the discussion of Australian politics, for example, shows that of the more than 26,000 users who participated during February to December 2011, the most active one per cent of users accounted for nearly two thirds of all tweets (the top ten per cent posted more than 90% of all #auspol tweets) – and that this leading group was considerably more likely to engage in @replying than the less active user groups. For other hashtags (such as widely televised, world events like #royalwedding), activity patterns are vastly different – here, the lead users account for a much smaller proportion of posts, and it is the 'long tail' which contributes the bulk of all tweets.

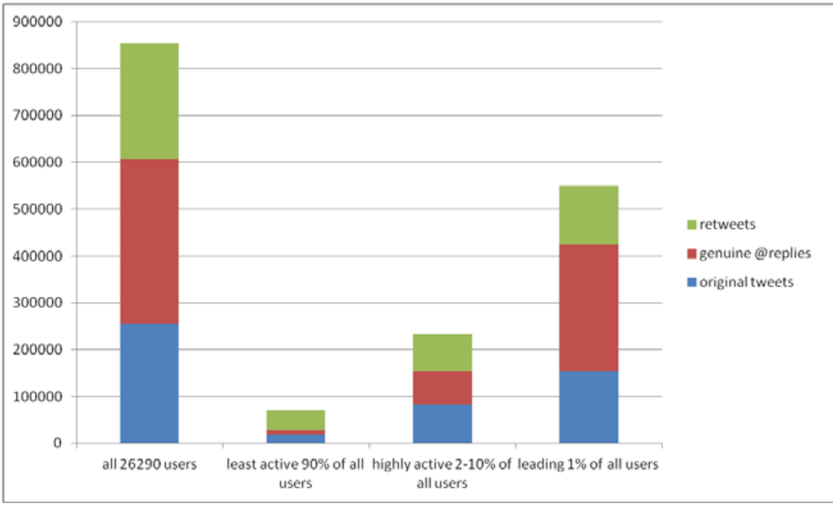

Fig. 3: contributions to #auspol made by the different groups of more or less active users (Feb. to Dec. 2011)

Again, the comparative work which is able to extend such analyses of individual hashtags to generate a more comprehensive view of how centralised or distributed individual Twitter events are, and how this correlates with the type of hashtag event in each case, depends crucially on the establishment of a standard set of metrics to describe these activity patterns. Such standardisation does not preclude hashtag-specific analysis, or aims to privilege the development of purely quantitative aggregate figures on hashtag usage over in-depth, qualitative study; rather, it serves as a crucial enabler for further qualitative research by pinpointing those leading users, key tweets, and other exceptional patterns which are most worthy of deeper analysis.

Beyond the Hashtag

The establishment of such standardised metrics for the study of Twitter interactions through hashtags enables new forms of comparative research which detects shared patterns and practices that transcend individual hashtags themselves. However, such work does not manage to overcome the fundamental limitations associated with hashtag-based approaches themselves: these necessarily cover only the tip of a communicative iceberg, and miss out, in particular, on a substantial amount of follow-on communication as users respond to hashtagged tweets but do not themselves include the hashtag in their @replies. More broadly, too, hashtag-based studies are appropriate only in communicative contexts where clearly established hashtags do exist – they are able to examine the particular form of political discourse which takes place in tweets carrying the #auspol hashtag, for example, but not the everyday political exchanges which take place, unhashtagged, right across the Australian Twittersphere.

Hashtag studies have been a popular tool for Twitter researchers in recent years not least because it is comparatively easy to capture a hashtag dataset, while the establishment of a representative or even comprehensive sample of general Twitter activity is considerably more difficult, especially for large populations of Twitter users (see e.g. Bruns & Liang, 2012): the former requires researchers to track just a single keyword, using readily available tools, while the latter must build on dedicated technology to identify and follow the public tweets of a potentially very large number of Twitter users on an ongoing basis. To date, few studies of Twitter populations at this level of comprehensiveness exist; future attempts to undertake them will have to wrestle especially with the prohibitive pricing regime for high-volume data access which Twitter has now established.

To the extent that they may be successfully carried out, such studies may again utilise the standard metrics outlined above, however. User metrics may be used, for example, to examine the distribution of diverse communicative approaches across a larger population of users, and could lead to the development of a systematic typology of Twitter users as described by their activity patterns (from users who specialise in posting original tweets only through to those who engage exclusively in retweeting the messages of others); as an aside, this could potentially also be used to automatically identify spambots and similar accounts with highly unusual tweeting patterns.

Tweet metrics, on the other hand, may be used on a population-wide basis to examine common diurnal patterns of Twitter activity (for example, to examine whether @replying or link-sharing take place more frequently at specific times of the day), or to highlight particular moments of heightened activity within the dataset. Where such analysis is possible in close to real time, it may enable the automatic detection of breaking news or crisis events, for example – similar to, but substantially extending beyond, the insights which Twitter's 'trending topics' already provide. Additionally, of course, tweet metrics may also be applied to the tweets sent by specific identified subsets of the overall Twitter population whose activities are being tracked; here, they generate insights which are comparable to those arising in hashtag studies, but may be able to transcend the inherent limitations introduced by focussing only on explicitly hashtagged tweets.

Finally, a more comprehensive study of Twitter activities amongst an identified population of users must also take into account more strongly the established follower/followee networks of these users. While studies proceeding from an analysis of a shared hashtag may assume that participating users are connected in the first place by their shared interest in the hashtag (which enables them to see one another's tweets even if they are not following each other), a population-wide study of Twitter patterns must build on the assumption that only the followers of a given user will be likely to see the tweets posted by that user.

This further complicates the analysis of such population-wide activity patterns; at the same time, however, the baseline patterns which a longitudinal study of Twitter use may be able to establish will also serve as an important point of comparison for the analysis of shorter-term hashtag events as we have outlined it above. Hashtag-based work alone may show the total volume of tweets responding to a certain event or issue, or may pinpoint certain users as leading contributors to the discussion; only in comparison to these baseline patterns, however, does it become possible for researchers to determine just how exceptional the hashtagged volume of tweets was, or how far from their standard patterns of interaction a user might have diverged in tweeting about a specific topic.

Conclusion and Reflections

In this paper, we have catalogued some recently developed, and potentially transferable methods and metrics for the study of public communication on Twitter; as a particularly prominent example of how social media platforms are remediating and transforming communication within the changing media ecology. In doing so we have demonstrated how a range of metrics and analytical techniques that address routine research questions in media and communication studies can help to make sense of the social media 'data deluge'.

However, there remain many new challenges for humanities and social science-inflected disciplines seeking to build on and extend data-driven approaches to internet communication. Two of the most significant of these concern methodology and disciplinary practices. First, media and communication researchers need to develop (and not just out-source) the appropriate technical skill and broader ‘code literacy’ sufficient to engage knowledgeably and critically with these methods – with broad consequences for the content and pedagogy of research training, and PhD programs especially. Second, there is much room for further development of multi-method approaches, integrating and innovating upon traditional qualitative methods (including close textul analysis and ethnographic approaches) in a ‘big data’ context, bearing in mind the critical ‘provocations’ for big data recently proposed by danah boyd and Kate Crawford (2011).

Beyond the practical methodological issues raised by the burgeoning field of data-driven media and communications research lie the political and pragmatic issues arising from competing regimes of data access, usage and control. For example, Twitter.com is effectively asserting monopoly rights on Twitter data through various technical and legal means, including the ban on web-based export of Twitter archives (makely the widely used archiving service Twapperkeeper ineffective for research purposes); and the choking off of access to its ‘firehose’ except via prohibitively expensive commercial providers such as Gnip, or by prior arrangement (as in the ‘gift’ of historical Twitter data to the Library of Congress). At the same time, the ‘open science’ and ‘open data’ movements propose a set of norms for scientific research that would ask us to make our original or processed datasets freely available for the use of our peers or the public (Rees, 2011) – creating a very complex set of problems for social science researchers who rely on third-party proprietary data such as Twitter archives.

With appropriate critical reflection, humanities and social science appropaches to the ‘scientific’ study of public communication, such as those discussed in this paper, may in fact offer a ‘special case’ of the politics of knowledge associated with the current turn to ‘big data’ and computational methods, *because* of their entanglement – even at the level of data collection – with the shifting business models of social media platforms, shifting and variable regulatory structures in relation to data access and use, as well as public anxieties around the control and use of our social data at a moment where ‘personal’ information and public communication are converging.

References

BERRY, DAVID, ED. (2012) *Understanding Digital Humanities.* London:Palgrave Macmillan.

BOYD, DANAH, SCOTT GOLDER & GILAD LOTAN (2010) Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter. Paper presented at 43rd Hawaii International Conference on System Sciences, Retrieved from http://doi.ieeecomputersociety.org/10.1109/HICSS.2010.412

BOYD, DANAH & KATE CRAWFORD (2011) 'Six Provocations for Big Data' *A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society,* Oxford.

BRUNS, AXEL. (2011, 12 Aug.). Twitter and the Royal Wedding, Pt. 1: Something Processed. Retrieved 5 July 2012 from http://mappingonlinepublics.net/2011/08/12/twitter-and-the-royal-wedding-pt-1-something-processed/

BRUNS, AXEL. (2005). *Gatewatching: Collaborative Online News Production*. New York: Peter Lang.

BRUNS, AXEL, & BURGESS, JEAN. (2012). The Use of Twitter Hashtags in the Formation of *Ad Hoc* Publics. Paper presented at the European Consortium for Political Research conference, Reykjavik, 25-27 Aug. 2011.

BRUNS, AXEL, & LIANG, YUXIAN EUGENE. (2012). Tools and Methods for Capturing Twitter Data during Natural Disasters. *First Monday* vol. 17, no. 4.

BRUNS, AXEL, & STIEGLITZ, STEFAN. (2012, under review). Quantitative Approaches to Comparing Communication Patterns on Twitter. Submitted to a special issue on "Methods for Analysing Social Media", *Journal of Technology in Human Services*.

CRAWFORD, KATE (2009) Following You: Disciplines of Listening in Social Media. Continuum: Journal of Media & Cultural Studies, 23(40): 525–35.

FULLER, MATTHEW, ED. (2008) *Software Studies: A Lexicon*. Cambridge, MA: MIT Press.

HERMIDA, ALFRED (2010) Twittering the News: The Emergence of Ambient Journalism. *Journalism Practice*, 4(3): 297-308.

MARWICK, ALICE & BOYD, DANAH. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media and* Society 13(1): 114-133

REES, JONATHAN (2011) Recommendations for independent scholarly publication of data sets. Creative Commons Working Paper, http://creativecommons.org

ROGERS, RICHARD (2009) The End of the Virtual. Inaugural address, University of Amsterdam. < http://www.govcom.org/rogers_oratie.pdf>