Anticipating Activity in Social Media Spikes

Desmond J. Higham
Department of Mathematics
and Statistics
University of Strathclyde,
26 Richmond Street,
Glasgow, G1 1XH, UK

Alexander V. Mantzaris
Department of Mathematics
and Statistics
University of Strathclyde,
26 Richmond Street,
Glasgow, G1 1XH, UK

Peter Grindrod
Mathematical Institute,
University of Oxford,
Oxford OX2 6GG, UK

Amanda Otley and **Peter Laflin**
Bloom Agency, Marshalls Mill,
Marshall Street, Leeds, LS11 9YJ, UK

Abstract

We propose a novel mathematical model for the activity of microbloggers during an external, event-driven spike. The model leads to a testable prediction of who would become most active if a spike were to take place. This type of insight into human behaviour has many applications, as it identifies key players who can be targeted with information in real time when the network is most receptive. The model takes account of the fact that dynamic interactions evolve over an underlying, static network that records "who listens to whom." Our fundamental assumption is that, in the case where the entire community has become aware of an external news event, a key driver of activity is the motivation to participate by responding to incoming messages. We validate the resulting algorithm on a large scale Twitter conversation concerning the appointment of a UK Premier League football club manager. We also find that the half-life of a spike in activity can be quantified in terms of the network size and the typical response rate.

Introduction

Digital footprints left by our online interactions provide a wealth of information for social scientists and present many new challenges in modelling and computation (Lazer et al. 2009). In addition to aiding our understanding of how humans interact and make decisions (Arak and Walker 2012), microblogging data offers the prospect of predicting future behaviour (Ciulla et al. 2012) and engaging in targeted intervention (Aral 2012). Commercial organisations, governments and charities are now able to interact with the general public during the course of an online, global conversation, and exploit opportunities to leverage current sentiment. We focus on the specific case where a rapid spike of activity can be attributed to an unpredictable external occurrence; examples considered here and in the Appendix include pivotal moments in a sporting event, breaking news, and climactic scenes in a TV programme. These dramatic, but short-lived, bursts of interest, with a typical half-life of between ten and twenty minutes, represent marketing opportunities for suitably agile players, as demonstrated by the cookie company Oreo, who produced an effective, and, subsequently award winning, tweet in response to a power failure during Super Bowl XLVII.

A number of authors have considered how information is passed in the setting of online social media. In (Cha et al. 2010) the number of followers, Retweets and mentions were used to quantify the influence of Twitter users, with the three measures yielding very different results. Similarly, (Kwak et al. 2010) ranked users by the number of followers and also by Google's PageRank algorithm. Related work in (Lerman, Ghosh, and Surachawala 2012) looked at how network structure affects dynamics of large scale information flow around news stories in Digg and Twitter. In (Centola 2010) the spread of behaviour was examined through artificially constructed, static, social interaction networks, with clustered-lattice structure found to be the most effective in terms of speed and reach. Dynamic analogues of the standard Katz centrality measures were tested on a large scale Twitter data set in (Laflin et al. 2013), and found to be compatible with the rankings produced by social media experts whose job is to identify key targets.

Our interdisciplinary work differs from previous studies in three main respects. First, rather than looking at the development of cascades *within* a community (for example the rise of a viral video) we focus on the setting where the relevant community has been roused by an *external* development. In this type of spike phase, because interest levels are high, there is a clear opportunity for targeted interventions to make an impact. Second, we develop a model that addresses both the dynamic nature of message-passing and the essentially static structure of the underlying "who listens to whom" network. Third, by making our key modelling assumption explicit and developing a simple algorithm that applies to large scale data sets, we produce a tool that can be employed in real time, predicting who will be the most active players as soon as a spike in volume has been detected.

Method and Results

In setting up a general modelling framework, we assume that no new associations are created during the short time scale of the spike; that is, we have a static underlying connectivity structure. To be concrete, we will discuss Twitter activity, but we note that the same principles apply to other time-dependent digital messaging systems. For the relevant set of N users, we let $A \in \mathbb{R}^{N\times N}$ denote a corresponding adjacency matrix where $a_{ij} = 1$ if user i is known to receive and take notice of messages from user j. Loosely, this might mean that i is known to be a Twitter follower of j, although in practice we have in mind the use of more concrete evidence that i cares about the tweets of j; for example via Retweets. For simplicity, we take a standard unit of time (one minute in the tests below) during which a user is assumed to send out at most one message. We let $s^{[k]} \in \mathbb{R}^N$ denote an indicator vector for the send activity at time k, so that $s_i^{[k]} = 1$ if user i tweeted in time interval k and $s_i^{[k]} = 0$ otherwise. Then simple bookkeeping tells us that

$$r^{[k]} = As^{[k]}, \tag{1}$$

where $r^{[k]} \in \mathbb{R}^N$ is such that $r_i^{[k]}$ counts how many messages were received by user i in this time interval.

We can now formalize our main modelling assumption. In words, the probability of a user tweeting at time $k + 1$ *is proportional to the number of significant tweets they have just received*, with proportionality constant denoted α, plus a basal rate. We therefore model $s^{[k]}$ as a discrete time Markov chain according to

$$\mathbb{P}\left(s_i^{[k+1]} = 1 \mid s^{[k]}\right) = b_i + \alpha\, r_i^{[k]}. \tag{2}$$

Here b_i denotes the basal tweet rate for user i and the second term on the right-hand side quantifies our assumption that, in the full attention span phase, activity is driven by a desire to join in with the current conversation and engage in topical "banter." Formally, a normalization factor should be included in the right-hand side of equation (2), to guarantee that probabilities lie between zero and one. However, we will see that for our purpose of ranking nodes, this is not necessary.

As general support for this key modelling assumption, we note that (Bakshy et al. 2012) found social influence to play a crucial role in the propagation of information on Facebook: "Those who are exposed [to friends' information] are significantly more likely to spread information and do so sooner than those who are not exposed." Similarly, (Wu et al. 2011) based on a Twitter study concluded that "although audience attention is highly concentrated on a minority of elite users, much of the information they produce reaches the masses indirectly via large population of intermediaries". Further empirical work appeared in (Lin et al. 2014), which looked at Twitter interactions under shared activity around eight major events during the 2012 U.S. presidential election. The study found that human behaviour changes during a "media activity," when information consumption is characterized by the availability of dual screening technology (television and hand held device) and real-time interaction. The authors proposed the term *media event-driven behavioural change* for this general effect, and showed that, for the data they collected, differences in behaviour were driven by the increasing attention given to a small cohort of elite users. Our work also focuses on this shared-attention, event-based setting, and the leadership role of central users, but considers behaviour when the whole network rapidly becomes aware of an item of breaking news.

From equations (1) and (2), the expected value $\mathbb{E}[s^{[k+1]}]$ evolves according to

$$\mathbb{E}[s^{[k+1]}] = b + \alpha A \mathbb{E}[s^{[k]}]; \tag{3}$$

see the Appendix. This type of iteration is familiar in many modelling and computation scenarios, and it is readily shown that as k increases $\mathbb{E}[s^{[k+1]}]$ generically lines up along a preferred direction that is independent of $s^{[0]}$; see the Appendix. If the spectral radius of A is below $1/\alpha$ then as $k \to \infty$ the resulting steady state value for $\mathbb{E}[s^{[k]}]$, which we denote by $s^\star$, satisfies

$$(I - \alpha A)s^\star = b. \tag{4}$$

If the time-scale for equation (3) to equilibrate is fast, relative to the time-scale of the activity spike, then $s^\star$ in equation (4) will summarize the state of the system. So, having constructed the matrix A and the right-hand side b from the current data, the vector $s^\star$ can be used to predict the relative activity level of each node in the event of a spike, a larger value of $s_i^\star$ suggesting that user i will be more active. In particular, the current top r components in the vector $s^\star$ give a prediction for who would be the r most active users if a spike were to erupt.

We also note that our new model can be used to explain the characteristic geometric decay in tweet volume observed in these examples following a peak of activity. In particular, the half-life of a spike can be predicted in terms of the typical community size; see the Appendix.

We may regard $s^\star$ as a *network centrality measure*; indeed it is related to the widely-used Katz centrality (Katz 1953), and can be interpreted independently from a combinatoric, graph-theoretic standpoint, as explained in the Appendix. In the special case where $\alpha = 0$, we do not make use of any underlying network information, and predict purely on the basal rate of each user. This provides a natural basis for testing the algorithm, and therefore validating our underlying hypothesis: does the use of $\alpha > 0$ add value to the prediction of who will be active during a spike? (We emphasize that, for simplicity, the single parameter α serves both to downweight long walks in the Katz centrality sense, and also to quantify the tendency of users to respond to incoming tweets.)

We address this question using a Twitter data set, with three further sets tested in the Appendix. In each case, we define a *business-as-usual* period where users operate at their basal rate and a *spike period*, where network activity has been dramatically increased by an external event. The basal tweet rate b_i for user i is taken to be their total number of business-as-usual period tweets. We also build the matrix A from business-as-usual data, setting $a_{ij} = 1$ if i received

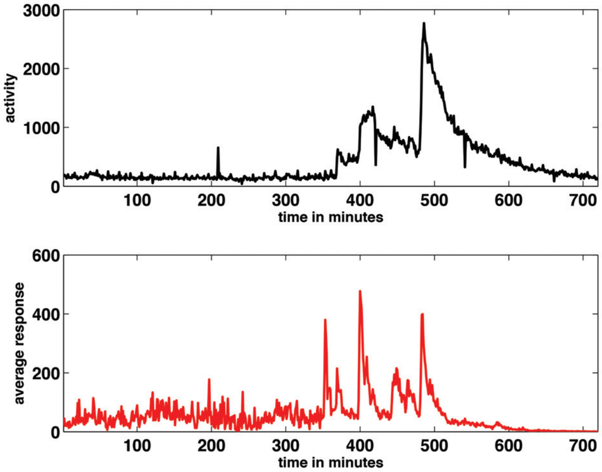

Figure 1: Upper: volume of tweets each minute for a conversation around Manchester United Football Club, May 9th, 2013. Lower: average responsiveness of the network, defined in equation (5).

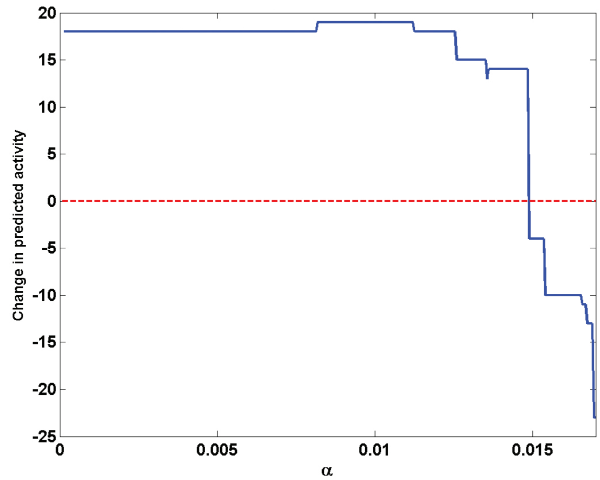

Figure 2: Solid line: change in activity of predicted top 100 tweeters during the spike phase in Figure 1, as a function of the response parameter, α. Dashed line: corresponding level for $\alpha = 0$.

at least one relevant tweet from j in this period, and setting $a_{ij} = 0$ otherwise.

Our experiment uses data collected on May 9th 2013 surrounding the appointment of David Moyes as manager of Manchester United Football Club, where a tweet is deemed to take part in the conversation if it contains one or more specified keywords. This consisted of 298,335 time-stamped directed message-passing events involving 148,918 distinct Twitter accounts. The upper picture in Figure 1 shows the volume of tweets each minute. The largest spike in volume, at 486 minutes, corresponds to the official announcement of Moyes' appointment. For the purposes of our test, we regard zero to 300 minutes as forming the business-as-usual period where users operate at their basal rate. We define the spike period as lasting from the peak time of 486 minutes to the time of 541 minutes at which the activity level has decayed by a factor of four.

As support for our modelling hypothesis that, in a spike phase, activity is driven by a desire to engage with incoming messages, we show in the lower picture of Figure 1 the *responsiveness* of the network, which we define as the number of tweets that a typical sender has seen in the previous one minute of their timeline. More precisely, we compute the *average responsiveness* over the kth one minute period as

$$\frac{1}{N_k} \sum_{p=1}^{N_k} \mathrm{rec}_k^{[p]}, \qquad (5)$$

where N_k denotes the number of tweets sent out in this minute and, for each such tweet, indexed by p, $\mathrm{rec}_k^{[p]}$ denotes the number of tweets that the sender received in the previous 60 seconds.

Now, we test the predictive power of the new measure $s^\star$ in equation (4) as a function of the response rate parameter, α. Figure 2 shows the change in *total spike period activity of the top 100 ranked users*, as a function of α. In other words, for each choice of α we use the business-as-usual information to compute $s^\star$, find the users with the 100 top-ranked values of $s^\star$ and then record the total number of tweets sent by this top 100 during the spike period. The figure shows the difference between the total activity of these users and those from the baseline value of $\alpha = 0$, corresponding to 560 tweets. As soon as α increases beyond machine precision level (around 10^{-16}), the top 100 list changes and the prediction improves, and this holds for a range of α values, relative to the $\alpha = 0$ case where no underlying connectivity is exploited. Comparing the fine details of $\alpha = 0$ with $\alpha = 0.01$, we find that both choices overlap strongly with the best possible top 100 (71 and 70 indices in common, respectively). However, the $\alpha = 0.01$ list benefits from including the 7th most active tweeter, who is missing from the $\alpha = 0$ list. We also remark that for this data set the ratio of the subdominant to the dominant eigenvalue of A has modulus of 0.9540, and hence over the time units of the spike period, since $0.9540^{54} \approx 0.08$, it is reasonable to assume that (3) has equilibrated. In the Appendix, the new algorithm is validated on data from three further Twitter conversations around high profile events.

In summary, our main aim was to put forward a novel mechanistic model for on-line human social interaction in the important setting where a community has been roused by an unpredictable, external influence. We couched our modeling assumptions in a very simple tractable model that allowed a concrete test to be performed on real Twitter data. It would be highy beneficial to develop rules of thumb, or theoretical arguments, to guide the choice of α, and of course there are many ways in which these ideas could be extended to more sophisticated models and subjected to further tests, notably at the micro-scale level, especially if larger relevant data sets become available. Finally, we emphasize that the testable hypothesis and resulting algorithm studied here are applicable in any digital social media setting where we pass information in real time to a pre-specified group of social neighbours.

Acknowledgements

We are grateful to Bloom Agency, Leeds, UK, for supplying the anonymised Twitter data, which has been placed in the public domain[1]. PG was supported by the Research Councils UK Digital Economy Programme via EPSRC grant EP/G065802/1 *The Horizon Digital Economy Hub*. DJH acknowledges support from a Royal Society Wolfson Award, a Royal Society/Leverhulme Senior Fellowship and an EPSRC/Digital Economy Fellowship. AVM was supported by the EPSRC and Bloom Agency via an Impact Acceleration Account secondment.

Appendix

Evolution of the Expected Value

Combining equations (1) and (2), we see that

$$\mathbb{P}\left(s_i^{[k+1]} = 1 \mid s^{[k]}\right) = b_i + \alpha \left(As^{[k]}\right)_i. \qquad (6)$$

We note that because $s_i^{[k+1]}$ takes only the value 0 or 1, its conditional expectation is simply the probability of taking the value 1, given $s^{[k]}$. So, taking conditional expectation in equation (6), we have

$$\mathbb{E}[s_i^{[k+1]}|s^{[k]}] = b_i + \alpha \left(As^{[k]}\right)_i.$$

Upon taking expected values, we then obtain

$$\mathbb{E}[s_i^{[k+1]}] = b_i + \alpha \left(A\mathbb{E}[s^{[k]}]\right)_i,$$

giving equation (3).

Stationary Iteration

For the iteration in equation (3), we have

$$\begin{aligned}\mathbb{E}[s^{[1]}] &= b + \alpha A \mathbb{E}[s^{[0]}] \\ \mathbb{E}[s^{[2]}] &= b + \alpha A \left(b + \alpha A \mathbb{E}[s^{[0]}]\right) \\ &= b + \alpha A b + (\alpha A)^2 \mathbb{E}[s^{[0]}],\end{aligned}$$

and the general pattern, which may be proved formally by induction, is

$$\mathbb{E}[s^{[n]}] = b + \alpha A b + \cdots + (\alpha A)^{n-1} b + (\alpha A)^n \mathbb{E}[s^{[0]}].$$

Under our assumption that $\alpha < 1/\rho(A)$, it follows that $\|(\alpha A)^n\| \to 0$ as $n \to \infty$, for any matrix norm $\| \cdot \|$. Hence, the influence of $s^{[0]}$ becomes negligible, and $\mathbb{E}[s^{[n]}]$ approaches

$$\sum_{i=0}^{\infty} (\alpha A)^i b,$$

which may be written $(I - \alpha A)^{-1} b$.

Katz-like parameter

In our notation, where A denotes an adjacency matrix, the kth power, A^k, has an (i, j) element that counts the number of directed walks from node i to node j. It follows that the infinite series

$$I + \alpha A + \alpha^2 A^2 + \cdots + \alpha^k A^k + \cdots$$

has (i, j) element that counts the total number of walks from from node i to node j of all lengths, where a walk of length k is scaled by α^k. Here "length" refers to the number of edges traversed during the walk. This series converges for $0 < \alpha < 1/\rho(A)$, whence it may be written $(I - \alpha A)^{-1}$.

The vector $c \in \mathbb{R}^N$ defined by $c = (I - \alpha A)^{-1}\mathbf{1}$, or, equivalently,

$$(I - \alpha A)c = \mathbf{1},$$

where $\mathbf{1} \in \mathbb{R}^N$ denotes the vector with all values equal to unity, therefore has ith element that counts the number of directed walks from node i to every node in the network, with a walk of length k scaled by α^k. This is one way to measure the "centrality" of node i, as first proposed by Katz (Katz 1953). In this way, α becomes the traditional attenuation parameter in the Katz setting, representing the probability that a message successfully traverses an edge. The measure $s^\star$ in equation (5) replaces the uniform vector $\mathbf{1}$ with b. Hence, the component $s_i^\star$ can be interpreted as a count of the total number of walks from node i to every node in the network, with walks to node j weighted by $b_j \alpha^k$. The introduction of b has therefore allowed us to weight the walk count according to basal dynamic activity.

Half-Life of a Spike

At the start of a spike, it is reasonable to suppose that $\mathbb{E}[s^{[0]}]$ in equation (3) is very large. We then have

$$\mathbb{E}[s^{[1]}] = b + \alpha A \mathbb{E}[s^{[0]}] \approx \alpha A \mathbb{E}[s^{[0]}],$$

and generally, in this spike phase,

$$\mathbb{E}[s^{[k]}] \approx (\alpha A)^k \mathbb{E}[s^{[0]}]. \qquad (7)$$

In the regime where $\alpha\rho(A) < 1$ it follows that the expected level of activity decays over time. More precisely, if we assume that the nonnegative matrix A is irreducible (that is, every node in the network has a path to every other) then the Perron–Frobenius Theorem (Higham 2008) says that there is a unique, real, positive, largest eigenvalue, λ_1 with corresponding nonnegative eigenvector v_1. We will expand $\mathbb{E}[s^{[0]}]$ as $\sum_{i=1}^{N} \beta_i v_i$, where $\{v_i\}_{i=1}^{N}$ are the eigenvectors of A, which we assume to span $\mathbb{R}^N$, with corresponding eigenvalues $\{v_i\}_{i=1}^{N}$ and with $\beta_1 > 0$. Then in equation (8),

$$\mathbb{E}[s^{[k]}] \approx \sum_{i=1}^{N} \beta_i (\alpha \lambda_i)^k v_i.$$

Since λ_1 is dominant, we have

$$\mathbb{E}[s^{[k]}] \approx \beta_1 (\alpha \lambda_1)^k v_1,$$

so

$$\mathbf{1}^T \mathbb{E}[s^{[k]}] \approx \beta_1 (\alpha \lambda_1)^k \mathbf{1}^T v_1.$$

[1] See the *Strathclyde MUFC Twitter Dataset* at http://www.mathstat.strath.ac.uk/outreach/twitter/mufc/index.php

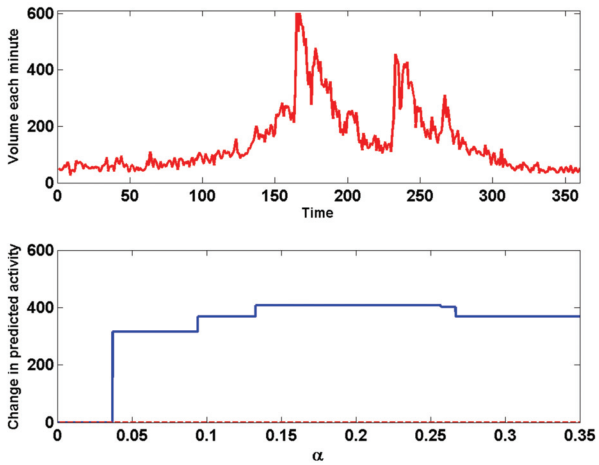

Figure 3: Upper: volume of tweets each minute for a conversation around a football match. Lower: change in activity of predicted top 10 tweeters during the spike phase, as a function of the response parameter, α. Dashed line: corresponds to $\alpha = 0$.

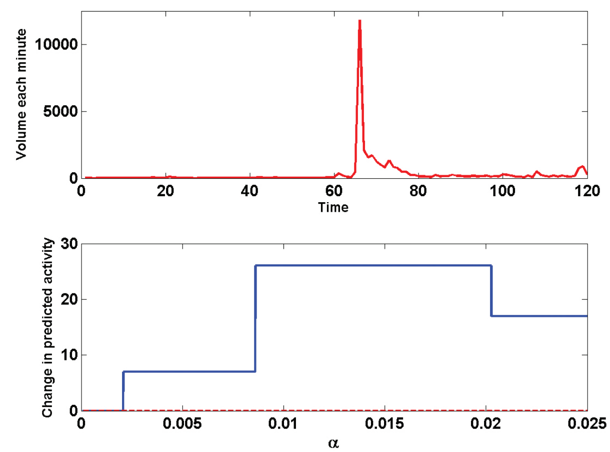

Figure 4: Upper: volume of tweets each minute for a conversation around a marketing event. Lower: change in activity of predicted top 10 tweeters during the spike phase, as a function of the response parameter, α. Dashed line corresponds to $\alpha = 0$.

We conclude that $\mathbf{1}^T\mathbb{E}[s^{[k]}]$, the overall expected network activity at time k, satisfies

$$\mathbf{1}^T\mathbb{E}[s^{[k]}] \approx C(\alpha\lambda_1)^k,$$

where C is a constant independent of k. The half-life then corresponds to $\widehat{k}$ time units, where

$$(\alpha\lambda_1)^{\widehat{k}} = \frac{1}{2},$$

leading to

$$\widehat{k} = \left| \frac{\log 2}{\log(\gamma)} \right|, \tag{8}$$

where γ is the product of the response rate α and the Perron–Frobenius eigenvalue of A, λ_1, which is bounded above by any subordinate matrix norm. Taking the standard $\| \cdot \|_1$ or $\| \cdot \|_\infty$ corresponds to forming the maximum in-degree or out-degree, respectively. So λ_1 may be regarded as roughly the maximum number of followers over all relevant users, and is hence a measure of community size.

Further Twitter Case Studies

Figure 3 presents results for a European football match: a Bundesliga encounter between Bayern Munich and Borussia Dortmund on May 4th, 2013. This involves 37,479 Twitter users. The upper picture shows Twitter volume per minute. We regard time zero to 130 minutes as the business-as-usual period, and define the spike period as starting at the peak of 165 minutes and finishing at 175 minutes, after which activity starts to increase again. This data is an order of magnitude smaller that the Manchester United data in Figure 1 of the main article. So we focused on the predicted top 10 users. The lower picture in Figure 3 shows the change in total spike period activity of this top ten as a function of α.

For Figure 4, we used data from a marketing event for the Yorkshire Tea Company on April 24th, 2013, where a range of tea lovers and celebrities, including Louis Tomlinson from pop band One Direction, took part in an Orient-Express style train journey around Yorkshire, UK, and were encouraged to publicize the event. In this case we have 9,163 Twitter users. The large spike at 66 minutes corresponds to awareness being raised about the presence of a One Direction member. We defined the business-as-usual period to last from zero to 65 minutes. The lower picture shows the change in spike activity of the predicted top ten as a function of the response parameter α.

Figure 5 shows results for a dual–screening conversation around an episode of the Channel Four UK television programme Utopia, involving 4,154 Twitter users. The spike at time 130 minutes corresponds to a particularly dramatic scene. We defined the spike to finish at 145 minutes, and took the business-as-usual period to last from time zero to 120 minutes. As before, the change in spike activity as a function of α is shown in the lower picture.

In each of these three further tests, we see that extra value is added by increasing α above zero; that is, by appropriately incorporating information about the underlying follower network that was built up in advance of the spike.

References

Arak, S., and Walker, D. 2012. Identifying influential and susceptible members of social networks. *Science* 337(6092):337–341.

Aral, S. 2012. Social science: Poked to vote. *Nature* 489(7415):212–214.

Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The role of social networks in information diffusion. In *Pro-*

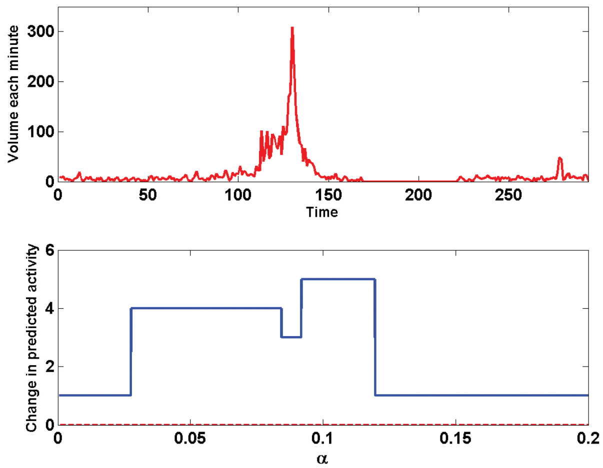

Figure 5: Upper: volume of tweets each minute for a conversation around a TV programme. Lower: change in activity of predicted top 10 tweeters during the spike phase, as a function of the response parameter, α. Dashed line corresponds to $\alpha = 0$.

ceedings of the 21st international conference on World Wide Web, 519–528. ACM.

Centola, D. 2010. The spread of behavior in an online social network experiment. *Science* 329(5996):1194–1197.

Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, P. K. 2010. Measuring user influence in Twitter: The million follower fallacy. *ICWSM* 10(10-17):30.

Ciulla, F.; Mocanu, D.; Baronchelli, A.; Gonçalves, B.; Perra, N.; and Vespignani, A. 2012. Beating the news using social media: the case study of American Idol. *EPJ Data Science* 1(1):1–11.

Higham, N. J. 2008. *Functions of Matrices: Theory and Computation*. SIAM Philadelphia.

Katz, L. 1953. A new status index derived from sociometric analysis. *Psychometrika* 18(1):39–43.

Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, 591–600. ACM.

Laflin, P.; Mantzaris, A. V.; Ainley, F.; Otley, A.; Grindrod, P.; and Higham, D. J. 2013. Discovering and validating influence in a dynamic online social network. *Social Network Analysis and Mining* 3(4):1311–1323.

Lazer, D.; Pentland, A.; Adamic, L.; Aral, S.; Barabási, A.; Brewer, D.; Christakis, N.; Contractor, N.; Fowler, J.; Gutmann, M.; et al. 2009. Computational social science. *Science* 323:721.

Lerman, K.; Ghosh, R.; and Surachawala, T. 2012. Social contagion: An empirical study of information spread on Digg and Twitter follower graphs. *arXiv preprint arXiv:1202.3162*.

Lin, Y.-R.; Keegan, B.; Margolin, D.; and Lazer, D. 2014. Rising tides or rising stars?: Dynamics of shared attention on Twitter during media events. *PLOS ONE* 9(5):e94093.

Wu, S.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Who says what to whom on twitter. In *Proceedings of the 20th international conference on World wide web*, 705–714. ACM.