


OPEN

# Unraveling the Origin of Social Bursts in Collective Attention

Manlio De Domenico<sup>1,2\*</sup> & Eduardo G. Altmann<sup>2,3</sup> 

**In the era of social media, every day billions of individuals produce content in socio-technical systems resulting in a deluge of information. However, human attention is a limited resource and it is increasingly challenging to consume the most suitable content for one's interests. In fact, the complex interplay between individual and social activities in social systems overwhelmed by information results in bursty activity of collective attention which are still poorly understood. Here, we tackle this challenge by analyzing the online activity of millions of users in a popular microblogging platform during exceptional events, from NBA Finals to the elections of Pope Francis and the discovery of gravitational waves. We observe extreme fluctuations in collective attention that we are able to characterize and explain by considering the co-occurrence of two fundamental factors: the heterogeneity of social interactions and the preferential attention towards influential users. Our findings demonstrate how combining simple mechanisms provides a route towards understanding complex social phenomena.**

The ability to filter the most relevant data out of a deluge of information is one of the characteristics of human intelligence. When this ability is coupled with individual's behavioral responses, like deciding to take an action based on the processed information, intriguing phenomena<sup>1</sup> such as collective attention might emerge. Like popularity, attention depends on a variety of both endogenous and exogenous factors that have effects on several aspects of human behavior, from timing patterns of activity<sup>2</sup> to peculiar responses to shocks<sup>3</sup>. The advent of social media and the possibility to record the simultaneous activity of millions of individuals allows the study of this type of phenomena on unprecedented large scales. In fact, such responses are often characterized by information cascades<sup>4–8</sup> and exhibit a rich dynamics with a long memory which is responsible, for instance, for the emergence of power-law distributed physical observables such as waiting times<sup>9,10</sup> and responses to social-media items<sup>11</sup>. This dynamics has been successfully modeled by a special class of self-exciting point processes known as Hawkes processes<sup>12</sup>, described by a self-reinforced dynamics where the likelihood of future events increases with the occurrence of a specific event.

Like online popularity<sup>13–19</sup>, collective attention is characterized by a quickly growing accumulated focus on a specific topic, e.g. presidential elections discussion on socio-technical systems, until a well identified peak of attention is reached, followed by a phase of decreasing interest with a slow decay<sup>20–22</sup>. The dynamical features of both the rise and decline of attention are still debated, although there is some evidence in support of power-law distributed activity<sup>9,10,14</sup> which is a signature of criticality in complex networked systems<sup>23</sup>. On the one hand, some studies succeeded in providing a description of collective attention dynamics while neglecting the effects of the underlying social structure<sup>24</sup>. In this case, the attention gathered by a content is understood as the result of an extrinsic factor – e.g., promotion of the content – acting upon two intrinsic factors, namely sensitivity to promotion and inherent virality<sup>25</sup>. On the other hand, recent studies highlighted the effects of the topological features, i.e. the underlying network of interactions, as well as of competing dynamics and memory time on the spreading phenomena observed in socio-technical systems<sup>26</sup>. Along this direction, many studies proposed different models based on the interplay between social structure and complex spreading dynamics to characterize the collective behavior observed in social media<sup>27</sup>, specially during special events such as the discovery of the “God particle”<sup>28</sup> or in response to real-world exogenous shocks such as disasters<sup>29</sup>. The interplay between system's topology and statistics of exogenous factors – such as news media – determine time-dependent network correlations that have been captured by more complex dynamical models of human activity, such as non-stationary<sup>30</sup> and non-linear<sup>31</sup> Hawkes processes and stochastic differential equations with Lévy noise<sup>32</sup>.

Here, we show that by combining two very simple mechanisms characterizing human activity it is possible to reproduce the most salient statistical features of extreme fluctuations<sup>33</sup> during collective attention in online

<sup>1</sup>CoMuNe Lab, Fondazione Bruno Kessler, Via Sommarive 18, 38123, Povo, TN, Italy. <sup>2</sup>Max Planck Institute for the Physics of Complex Systems, 01187, Dresden, Germany. <sup>3</sup>School of Mathematics and Statistics, The University of Sydney, 2006, NSW, Australia. \*email: [mdedomenico@fbk.eu](mailto:mdedomenico@fbk.eu)

Event	Interactions	Users	From	To	Days	Keywords
Boston Attack	9,480,331	4,377,184	2013-04-15	2013-04-22	7.0	"boston", "bomb", "BostonMarathon"
Papal Conclave (Pope Francis)	5,969,189	2,064,866	2013-02-25	2013-03-19	22.0	"pope", "benedict", "pontifex", "resign", "conclave", "vatican"
Paris Attacks	4,163,947	1,896,221	2015-11-13	2015-11-15	2.0	"#Paris" (search and streaming), "#Parigi" (streaming only)
NBA Finals	2,150,187	747,937	2015-06-09	2015-06-21	12.0	"#nbafinals"
UEFA Champions League Final	1,673,492	677,145	2016-05-27	2016-06-01	5.0	"#UCLfinal", "#RealAtletico", "#Champions"
Cannes Film Festival	1,180,173	438,537	2013-05-06	2013-06-03	28.0	"cannes film festival", "cannes", "#cannes2013", "#festivalcannes", "#palmdor", "canneslive"
Gravitational Waves Discovery	721,590	362,086	2016-02-10	2016-02-16	6.0	"ligo", "#gravitationalwaves", "#ligo", "gravitational waves", "#gravitational waves", "gravitational #waves", "onde gravitazionali", "#OndesGravitacionnelles", "Ondas gravitacionales", "Ondes Gravitacionnelles", "#ondas #gravitacionales", "#ondas gravitacionales"
Sanremo Italian Music Festival	461,838	56,562	2016-02-13	2016-02-13	1.0	"sanremo"
50th Anniv. of M.L. King's "I have a dream" speech	398,230	327,707	2013-08-25	2013-09-02	8.0	"Martin Luther King", "#ihaveadream"

**Table 1.** Information about network data sets used in this study. Note that data for *Cannes Film Festival* and *50th Anniv. of M.L. King's "I have a dream" speech* is a subset of the data used in Omodei *et al.*<sup>22</sup>.

social systems, without focusing on the evolution of the underlying dynamics. More specifically, we consider a preferential attachment process, related to individual's neighborhood and social connectivity that characterizes the network topology, and a preferential attention process, a cognitive dynamics related to individual's attention bias towards specific users of the network.

## Results

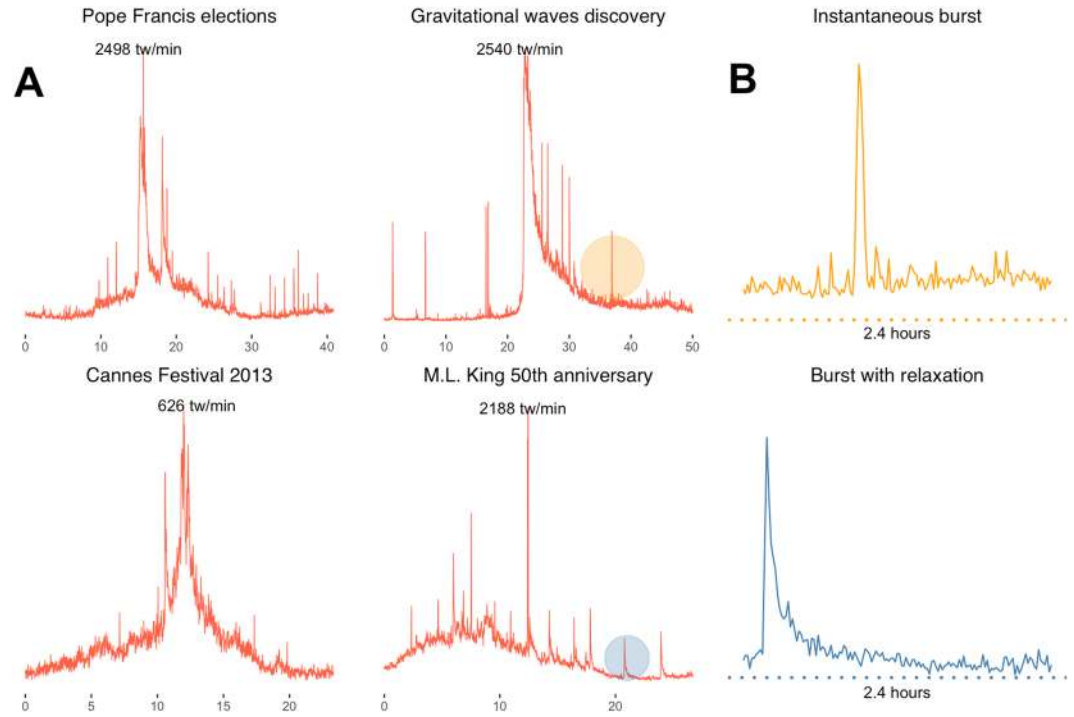
**Overview of the data sets.** In this work, we analyze the online activity of millions of users posting millions of messages in Twitter, a popular microblogging platform, during nine special events. We focus on events of wide public interest spanning different topics, such as the elections of Pope Francis (religion), NBA finals (sport), the discovery of gravitational waves (science), and the Cannes Film Festival (culture). The data sets consist of the second-by-second online activity that for the subsequent analysis has been aggregated at the time scale of  $T = 1$  min. More details about the data are provided in the Materials Methods section below, in particular in Table 1.

**Analysis of bursty activity due to collective attention.** Figure 1A shows the temporal evolution of the collective attention gathered by four special events. A striking feature, observed in all events regardless of their type (e.g., political, religious, cultural, scientific), is the bursty behavior of the social system: spikes of activity appear to be randomly placed on top of a more smooth temporal variation. Figure 1B shows that the spikes are extremely sharp in time, characterized by an abrupt increase followed by either a decrease of activity within one time unit (1 minute, in the figure) or by a slightly slower decrease of activity resembling the relaxation of a system's response to some stimulus. The main goal of this manuscript is to provide a statistical characterization of these spikes (or bursts of activity) and to discuss possible mechanisms that account for them.

A first insight on the origin of the spikes is obtained by decomposing the overall activity into its components due to individual's lone activity ("Tweet") – posting messages related to the event which do not involve other users – and to social interactions, such as endorsing ("Retweet") or replying to ("Reply") other individual's posts. Figure 2 shows that bursts dominated by both individual and social activities exist. Counting the contribution of each activity to many different bursts, compared against random expectations, reveals that the social interactions (retweets and replies) are more frequently responsible for the spikes (see Suppl. Figs. 1–2).

**Characterizing bursty activity in collective attention.** What type of mechanisms can be responsible for the spiky online activity summarized above? Recent studies attempted to relate the overall collective activity to peculiar characteristics of the underlying social structure or the influence of endogenous and exogenous factors<sup>31</sup>. The extremely fast and socially-dominated nature of spikes point towards a mechanism of reinforcement of collective behaviour taking place endogenously in the social network. Our hypothesis is that the variety of fluctuations observed in empirical data are due to the interplay between topological effects, related to the individual's neighborhood and social connectivity, and cognitive effects, related to the individual's bias towards activity from specific users. Both effects are known to concentrate the attention in the few most connected users. This motivates us to search for mathematical models that account for the spiky collective attention observed in online platforms such as Twitter and that just depend on individual's relationships and interactions.

We concentrate on the typical case of spikes generated by social activities in response to previous messages. Once a message  $i$  is posted, the  $k_i$  followers of the source user who posted it can act socially (i.e., in Twitter this might correspond to a Reply or a Retweet, corresponding to a direct comment or an endorsement, respectively). In our simple model, the multiple factors affecting this response are reduced to two:  $p_A(t)$  the probability of a



**Figure 1.** Social bursts of collective attention during exceptional events. **(A)** Volume of activity in tweets/minute (y-axis) as a function of time (x-axis, measured in hours) observed in the microblogging platform Twitter and measured during special events (Pope Francis’ election in 2013, the discovery of gravitational waves in 2016, the Cannes Film Festival in 2013, and the 50th anniversary of Martin Luther King’s most famous speech in 2013). **(B)** Bursts decay either instantaneously (top) or with some characteristic relaxation dynamics (bottom). The collective activity shown here aggregates the number of messages and the social actions they trigger:  $N(t) + R_{retweets}(t) + R_{replies}(t)$ .

follower being active and  $q_i(t)$  the probability of an active follower to react. The extremely short time scales of the spikes suggests that the reactions to a message are dominated by the immediate followers of the source user, instead of long/deep cascades of interactions in the network. With this simplifying assumptions, the probability that the message  $i$  triggers  $R_i$  social activities (responses) at time  $t$  is given by

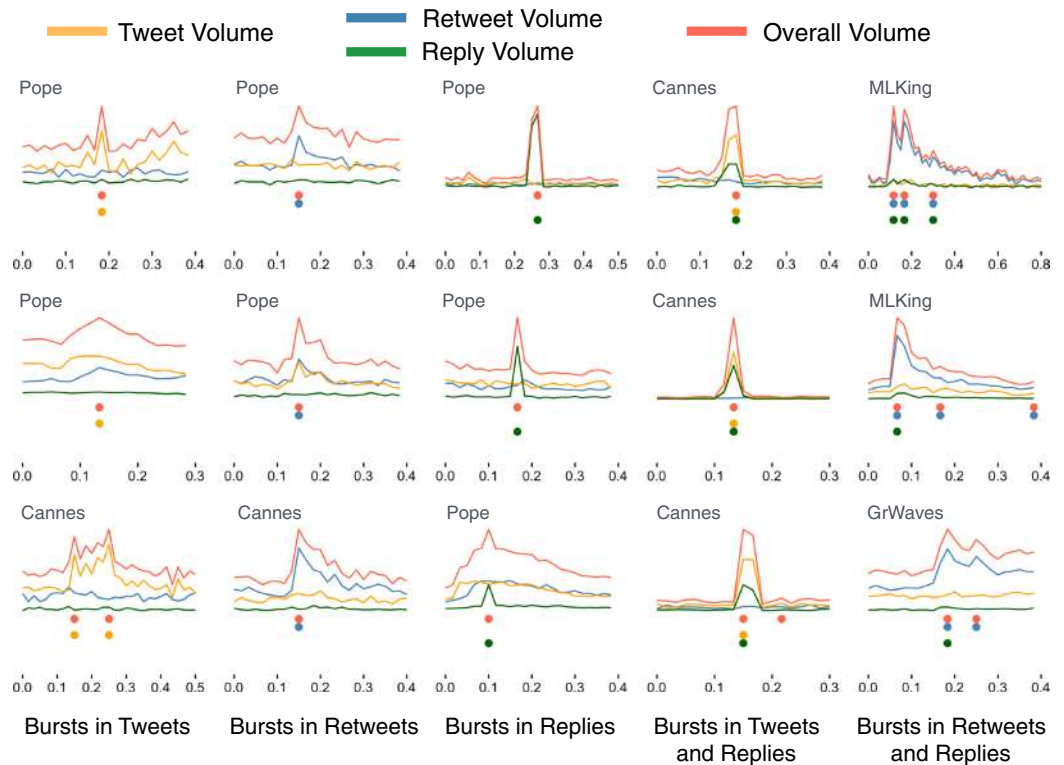
$$P(R_i(t)|i) = B(k_i p_A(t), q_i(t)), \tag{1}$$

where  $B$  is the Binomial distribution with  $k_i p_A(t)$  samples and probability  $q_i(t)$ . The overall social activities  $R(t)$  at time  $t$  is obtained summing the number of triggered responses  $R_i(t)$  over all  $N(t)$  messages contributing to social activities at time  $t$  as

$$R(t) = \sum_{i=1}^{N(t)} R_i(t) \approx p_A(t) \sum_{i=1}^{N(t)} q_i(t) k_i, \tag{2}$$

where the approximation is based on the expected number of interactions  $k_i p_A(t) q_i(t)$ , the average of distribution in Eq. (1). In Eq. (2), we consider messages to be randomly placed in the network so that for each message the user associated to  $i$  (with  $k_i$  and  $q_i$ ) is randomly chosen. In particular, we consider  $k_i$  to be a random sample of the degree distribution of the network  $\rho(k)$ . Our analysis of empirical data reveals that the duration of bursts due to social actions is, on average, shorter than 5 minutes, with 15 minutes as an upper bound (see Suppl. Fig. 2). Due to this extremely short time scales of the duration of the bursts, and similar short time scales for the social reactions to posted messages (see Suppl. Fig. 3), in our model we estimate  $N(t)$  (the number of messages contributing to social activities at time  $t$ ) simply as the average number of messages published in a window of time around  $t$  (see ref. <sup>11</sup> for a more detailed account of the slow temporal decay of the number of social interactions to a message). It is worth remarking that messages should not be necessarily produced at time  $t$ , but they can be posted before without triggering social interactions before time  $t$ .

Equation (2) defines our simple model for collective attention, and different scenarios are obtained by specifying the network (its degree distribution  $\rho(k)$ )  $p_A(t)$ , and  $q_i(t)$ . The probability of a user to be active  $p_A(t)$  simply re-scales the number  $R(t)$  of social activities and will thus not be relevant in our explanation of the spikes. The two critical parameters in the different scenarios are  $k_i$  – the number of users that receive the message  $i$  – and  $q_i(t)$  – the probability of user  $i$  to act socially (retweet/reply). We consider three different scenarios of increasing complexity:



**Figure 2.** Demultiplexing collective attention into specific activities. Different social actions contribute to online collective attention. Here we disentangle three different actions – lines and dots in different colors in each panel – and show their intensity ( $y$ -axis) over time ( $x$ -axis), as well as their combination (Overall Volume). Each of the 15 panels shows a spike reflecting a burst of activity (the time of the spike is indicated by a red dot). Spikes were automatically detected in the time series of overall volume (see Materials and Methods section for details on the detection method). Each column of three panels shows spikes due to different social actions, as indicated by the label in the lower corner (the first column shows three spikes originated from tweets, the second column shows three spikes originated from retweets, etc.). Multiple spikes occur during different exceptional events, the event of each spike is indicated in the label in the top-left corner of each panel.

1. Homogeneous:  $q_i(t) = q(t)$  is independent of  $i$  and  $\rho(k)$  is sharply peaked around an average degree  $\langle k \rangle$  (e.g.,  $\rho(k) \sim \text{Pois}(\langle k \rangle)$ ). In this case, the role of  $q(t)$  is to simply re-scale  $p_A(t)$ , which are both assumed to have a smooth temporal dependence not related to the spikes. Fluctuations in this scenario are expected to be small because of the well-behaved degree distribution  $\rho(k)$ , so that this scenario acts as a null model.
2. Heterogenous: we incorporate to the previous scenario the well-known fact that  $\rho(k)$  is a fat-tailed distribution, decaying as  $\rho(k) \sim k^{-(1+\mu)}$  for  $k \gg 1$ . Typically  $1 < \mu < 2$  and in the specific case of Twitter,  $\mu \simeq 1.2$  was measured<sup>34</sup>. Much larger fluctuations are expected in this scenario because of the strong variations in  $k_i$  for different  $i$ , i.e., the messages coming from hubs ( $k_i \gg \langle k \rangle$ ) are expected to receive much more interactions than messages from typical nodes ( $k_i \approx \langle k \rangle$ ).
3. Preferential attention: we incorporate to the previous scenario the fact that reaction to a message is more likely if it comes from a user that is perceived as important or central. The simplest proxy for such an importance is the degree of the message creator and thus we use  $q_i(t) \propto k_i$ .

For each of the scenarios, the sum in Eq. (2) effectively considers samples of distribution with short (case 1) or fat (case 2 and 3) tails. The restriction  $1 < \mu < 2$ , valid for all degree distributions  $\rho(k) \sim k^{-(1+\mu)}$ , ensures that  $\langle k \rangle$  is well defined in scenario 2. In contrast, scenario 3 effectively corresponds to drawing samples from a distribution with diverging mean because  $q(k)\rho(k) \sim k^{-\mu}$  (i.e., the exponent is reduced by one due to  $q \sim k$ ). See Materials and Methods, Model with preferential attention, for further details.

**Revealing the mechanisms behind collective attention.** The mechanisms behind collective attention can be revealed by testing to which extent the scenarios above describe the observations. We are interested in the spikes observed in the data, an extreme case of variability of the activity. Here, the data is represented by a time series of length  $L$  encoding the collective activity of the social network over time. We divide this time series into non-overlapping windows of size  $\ell$  and, for each window  $w = 1, 2, \dots, L/\ell$ , we quantify the spikiness  $S_w$  in the window as the ratio between the maximum and the mean volume  $R(t)$  of social responses, for  $t$  in the window:

$$S_w = \frac{\max_{t \in w} R(t)}{\langle R(t) \rangle_{t \in w}}. \quad (3)$$

The overall number of posted messages in each window is indicated by  $N_{T,w} = \sum_{t \in w} N(t)$  and we consider  $N(t)$  in Eq. (2). First we discuss the expectations for the dependence of  $S_w$  on  $N_{T,w}$  for the three scenarios of our model. As argued above, the scenarios correspond to random sampling of three fundamentally different types of distribution (short, heavy, and extremely heavy-tails). Accordingly, the  $\mathbb{E}[R]$  and  $R^{\text{Max}}$  – the expected value of  $R(t)$  and its largest value in  $\ell$  independent realizations, respectively – scale differently with the number of messages  $N_{T,w}$ , leading to the following estimations of the spikiness  $S_w$  (see Materials and Methods, Sum of fat-tailed variables and Maxima):

1. Homogeneous:  $S_w \sim 1/\sqrt{N_{T,w}}$ , i.e., the usual central-limit-theorem decay (i.e., spikiness is not expected for large values of  $N_{T,w}$ )
2. Heterogeneous:  $S_w \sim 1/N_{T,w}^{1-1/\mu}$ , i.e., a slower decay of  $S_w$  (i.e., spikiness persist for larger values of  $N_{T,w}$ )
3. Preferential attention:  $S_w$  does not depend on  $N_{T,w}$  or, at most, decays slower than algebraic.

The scaling (“ $\sim$ ”) relationships above hold for  $N_{T,w}(t) \gg 1$ , the usual setting of the generalized central limit theorem (see Materials and Methods). When  $N_{T,w}(t) \approx 1$ ,  $R(t)$  will follow the distribution of  $q_i(t)k_i$ . As anticipated, the activation probability  $p_A(t)$  just rescales the triggered social activities  $R(t)$  in Eq. (2) and therefore it cancels out in the ratio defining  $S_w$  in Eq. (3).

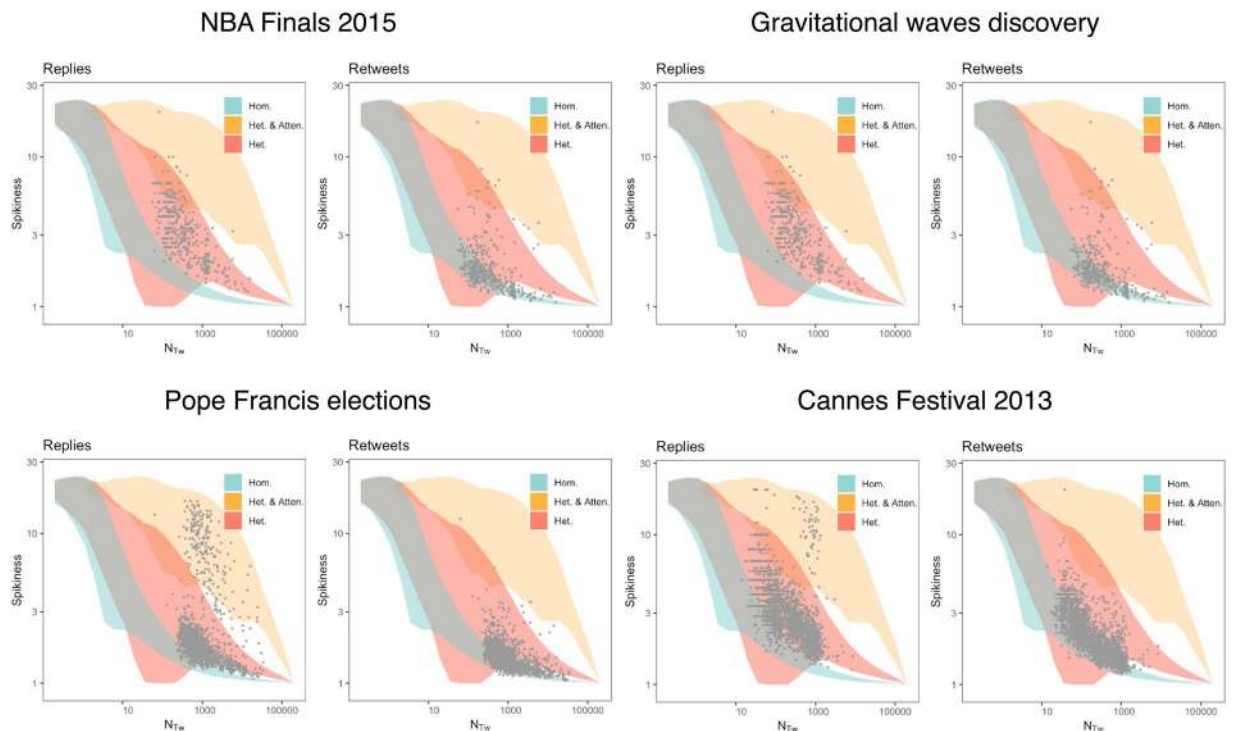
In the analysis of the empirical data, typical choices for  $\ell$  range from 20 minutes to a few hours: it can not be too small or too large, to allow for a significant number of samples to be analyzed. Each time window consists of  $\ell$  measurements, because we have built the data sets at 1 minute resolution.

In order to allow for a meaningful comparison between the data and the results obtained from the model, we generate several independent Monte Carlo realizations of the three scenarios of our model. In each realization, we generate the overall collective activity – including posting messages and social responses – in a window of size  $\ell$  and compute the corresponding spikiness  $S$  as in the empirical case reported above. This is done for increasing values of  $N_{T,w}$  and using  $N(t) = N_{T,w}/\ell$ . This choice of  $N(t)$  is justified by the short time scales of the reactions to tweets – as argued after Eq. (2) – and can be viewed as a lower bound on the number of messages actively generating reactions at time  $t$ . The value of the spikiness expected from the models and its corresponding variability are calculated over the ensemble of Monte Carlo realizations. For each scenario, we build the 90% confidence interval around the expected spikiness and we evaluate whether the pairs  $(N_{T,w}, S_w)$  measured from the data lie within this region. The results are shown in Fig. 3. The fluctuation analysis reveals some remarkable features of collective attention. The three scenarios introduced in this work account, all together, for the observations. Most empirical points show small spikiness  $S$  and are well explained by the homogeneous and heterogeneous scenarios. These correspond to the parts of the time series with small fluctuations (notice that Fig. 3 includes  $S$  for all windows  $w$ , not only the spikes reported in Figs. 1 and 2). From the point of view of our model, this means that the preferential attention mechanism does not play a role for these tweets (or during these periods). However, possibly the most interesting observation of Fig. 3 is that many points lie outside the range of the heterogeneous model and can only be understood taking the preferential attention mechanism into account. These points correspond to the most pronounced peaks reported in Figs. 1 and 2, the ones that we are mostly interested in this paper. This suggests that the skewed degree distribution of online social network is not enough to explain the spikes and that the network asymmetry is further amplified by additional mechanisms (such as the preferential attention considered in our model). Figure 3 reveals also that the statistics of social activities can vary widely within the same event, as in the case of Pope Francis election, where replies fluctuations are well explained either by the preferential attachment or the preferential attention models. In general, the spikes of Retweets are compatible with the heterogeneous model, while the spikes of Replies are larger than expected by the heterogeneous model and can only be accounted in the preferential attention scenario.

## Discussion

During events of special relevance, collective activity is usually more frenetic – i.e. the probability of posting is sufficiently high to guarantee a larger number of messages posted to the social system, typically well above 50 messages per minute – and the overall interest in the subject is driven by external factors. On top of this (smooth) overall tendency, extremely large fluctuations can be observed in form of spikes. These spikes have very short duration (often less than a 1 minute) and reflect a burst of activity and a dramatic concentration of the total social attention. The existence of these sharp spikes is the main empirical finding of our manuscript, which we further investigated through data analyses and comparisons to simple models. We find that spikes can have different origins but that most of them are due to social activities – such as Replies or Retweets in Twitter – in response to messages coming from well connected nodes.

In order to understand the origin of the spikes, we introduced a simple stochastic model that captures known effects that lead to extreme fluctuations such as the ones observed in social bursts of collective attention. It incorporates two fundamental mechanisms: the preferential attachment process, related to individual’s neighborhood and social connectivity that characterize the observed network topology, and a preferential attention process, a cognitive dynamics related to individual’s attention bias towards specific users of the network. In this work we considered an heterogeneous connectivity distribution scaling as  $k^{-2.2}$ , according to independent measurements<sup>34</sup>, and attention bias linearly proportional to the connectivity  $k$ . Comparing the model predictions with Twitter data, we find that the more extreme bursts of collective behavior – typically in form of Replies – are larger than what could be expected from the preferential attachment process alone and can only be understood



**Figure 3.** Fluctuation analysis of social bursts during collective attention. Spikiness  $S$  – Eq. (3) – is plotted against number of tweets –  $N_{T,w}$  – for two social activities (replies and retweets) during four exceptional events. Each dot is the result (empirical data) obtained in a time window  $w$  of size  $\ell = 20$  minutes (i.e.,  $N_{T,w} = 1000$  indicates an average of 50 posts per minute). Shaded areas indicate the 90% confidence around the expected  $S$  obtained simulating our model in the three scenarios: (i) homogenous social structure with uniformly distributed attention (“Hom.”); (ii) social structure obtained from preferential attachment with uniformly distributed attention (“Het.”); (iii) social structure obtained from preferential attachment with preferential attention (“Het. & Atten.”).

considering that the preferential attention process further amplifies the skewness of attention towards specific contents.

Our results show that two simple mechanisms are able to reproduce the statistical features of the appearance of spikes during exceptional events and our approach provides a procedure to measure the existence and the influence of preferential attention during events triggering collective attention.

## Materials and Methods

**Overview of the data sets.** Messages posted by users in Twitter, a popular microblogging platform, have been collected using the streaming real-time provided by Twitter API platform, filtered by the specific keywords reported in Table 1. By default, Twitter limits to 1% of the overall number of messages per second that can be retrieved from the streaming API. However, when the fraction of tweets concerning specific keywords is smaller than 1% of the global volume, Twitter does not apply limitations and the complete flow of information is collected. When this is not the case, Twitter provides messages of warning, reporting the cumulative number of missed tweets. For all events considered in this work, the estimated completeness of the sample is above 95%. Because of Twitter policies, the data sets (original tweet IDs) are available upon request. The main network data (adequately anonymized) analyzed in this work, are made publicly available at this: <https://github.com/manlius/SocialBursts>.

**Burst detection algorithm.** To determine the presence and the temporal position of bursts we first identify all local maxima (or peaks) in the time series of the overall tweets volume (i.e., including all tweets with and without social actions). In this stage, we used the *findpeaks* function of the R package *pracma*, publicly available at the: <https://cran.r-project.org/web/packages/pracma/index.html>. As reported in the corresponding documentation, this function is quite general as it relies on regular patterns to determine where a peak is located, from beginning to end. It is used with default values, except for the parameter threshold which controls the minimum value a peak should have to be considered as such. For each exceptional event considered in this study, we set this parameter to the corresponding observed median of the overall volume of tweets across time. For each identified peak, the output of this stage includes also the time at which a peak  $p$  started ( $t_{start}^{(p)}$ ) and ended ( $t_{stop}^{(p)}$ ) (i.e., the temporal range).

In a second stage, the temporal range  $[t_{start}^{(p)}, t_{stop}^{(p)}]$  is used to determine the peak duration  $\Delta t^{(p)} = t_{stop}^{(p)} - t_{start}^{(p)}$  and the volume of activities for each action (tweet without social actions, replies, and retweets) separately, as well as their overall volume:

$$\begin{aligned}
 N^{(p)} &= \sum_{t=t_{start}^{(p)}}^{t_{stop}^{(p)}} N^{(p)}(t) \\
 R_{retweets}^{(p)} &= \sum_{t=t_{start}^{(p)}}^{t_{stop}^{(p)}} R_{retweets}^{(p)}(t) \\
 R_{replies}^{(p)} &= \sum_{t=t_{start}^{(p)}}^{t_{stop}^{(p)}} R_{replies}^{(p)}(t) \\
 V^p &= N^{(p)} + R_{retweets}^{(p)} + R_{replies}^{(p)}.
 \end{aligned}$$

For each peak, the relative contribution of each action with respect to its overall volume is calculated:

$$\begin{aligned}
 r_{tweets}^{(p)} &= N^{(p)} / V^{(p)} \\
 r_{retweets}^{(p)} &= R_{retweets}^{(p)} / V^{(p)} \\
 r_{replies}^{(p)} &= R_{replies}^{(p)} / V^{(p)}.
 \end{aligned} \tag{4}$$

Similarly, the probability that an observed activity is a tweet, a retweet or a reply, is estimated as follows:

$$\begin{aligned}
 p_{tweets} &= \sum_t N(t) / \sum_t V(t) \\
 p_{retweets} &= \sum_t R_{retweets}(t) / \sum_t V(t) \\
 p_{replies} &= \sum_t R_{replies}(t) / \sum_t V(t),
 \end{aligned} \tag{5}$$

where the sum is over the whole time course of the considered exceptional event.

In a third stage, for each peak and each action  $i = \{\text{tweets, retweets, replies}\}$  we assume a binomial null model with mean  $\mu = V^{(p)}p_i$  and variance  $\sigma^2 = V^{(p)}p_i(1 - p_i)$  and check if the inequality  $V^{(p)}r_i^{(p)} > \mu + S\sigma$  is satisfied. We use  $S = 2.6$  for all statistical tests, corresponding to require the volume of peaks to be in the top 1% of the distribution or, equivalently, 2.6 standard deviations above the population mean. The peaks passing our tests are classified as *social bursts* and were used in the computations leading to Fig. 2 above and to Figs. 1 and 2 of the Supplementary Material.

**Model with preferential attention.** We consider that the probability of a social action is not a constant  $q$  but depends on the message  $i$  as  $q_i = qk_i^\alpha$ . From Eq. (2) we obtain

$$R_N = p_A \sum_{i=1}^N q_i k_i \propto \sum_{i=1}^N k_i^{1+\alpha}.$$

This can be viewed as a sum of  $N$  i.i.d. random variables  $x_i \equiv k_i^{1+\alpha}$  which then, a part from pre-factors, is in the form of Eq. (6). The distribution of  $x_i$ ,  $\rho(x)$ , is related to  $\rho(k)$  by  $\rho(x) = \rho(k) / dx/dk$ . For the power law case  $\rho(k) \sim k^{-(1+\mu)}$  we obtain

$$\rho(x) \sim k^{-(1+\mu+\alpha)} \sim x^{-(1+\mu+\alpha)/(1+\alpha)}.$$

This means that the model with preferential attention is equivalent to the model with heterogeneous networks with a modified exponent, obtained by the mapping of the exponents  $1 + \mu \mapsto 1 + \mu' = (1 + \mu + \alpha)/(1 + \alpha)$ . In the text we consider the case  $\alpha = 1$  which leads to a modified exponent  $\mu' = \mu/2$ . In particular,  $1 < \mu < 2$  is mapped to  $0.5 < \mu' < 1$ . The degree distribution of networks is  $\mu > 1$  and therefore an effective exponent with  $\mu < 1$  (the third case discussed below,  $0 < \mu < 1$ ) can only be achieved through the incorporation of preferential attention.

**Sum of fat-tailed variables.** Let  $x \geq 0$  be a random variable with distribution  $\rho(x)$  such that  $\rho(x) \sim x^{-(\mu+1)}$  for large  $x$ , with  $\mu > 0$  (fat tails). We are interested in the sum of  $N$  independent samples of  $x$

$$R_N = \sum_{i=1}^N x_i. \tag{6}$$

Following ref. <sup>33</sup>, the following cases can be described:

1.  $\mu \geq 2$  In this case, which includes also distributions  $\rho(x)$  with short tails (such as the Poisson distribution in scenario 1.), both moments  $\langle x \rangle$  and  $\langle x^2 \rangle$  exist and for large  $N$  the usual central limit theorem applies such that  $\mathbb{E}[R_N] = \langle x \rangle N$  and variance  $\mathbb{V}[R_N] = 2\sigma_x^2 N$ . Therefore, fluctuations are small and decay with  $N$  as

$$\frac{\sqrt{\mathbb{V}[R_N]}}{\mathbb{E}[R_N]} \sim \frac{1}{\sqrt{N}}.$$

2.  $1 < \mu < 2$  In this case,  $\langle x \rangle$  exists but  $\langle x^2 \rangle$  does not. The expected value  $\mathbb{E}[R_N] = \langle x \rangle N$  holds, but the fluctuations increase dramatically. In particular,  $\mathbb{V}[R_N]$  diverges with  $N$  as

$$(R_N - \mathbb{E}[R_N])^2 \sim N^{2/\mu}, \quad (7)$$

and therefore

$$\frac{\sqrt{(R_N - \mathbb{E}[R_N])^2}}{\mathbb{E}(R_N)} \sim N^{(1-\mu)/\mu}. \quad (8)$$

This still decays to zero because  $(1 - \mu)/\mu < 0$ .

3.  $0 < \mu < 1$  In this case  $\langle x \rangle$  is not defined and

$$\mathbb{E}[R_N] \sim N^{1/\mu}. \quad (9)$$

**Maxima.** Our measure of spikiness  $S_w$  defined in Eq. (3) consider the block- $\ell$  maximum of  $R$ , denoted by  $R^{\text{Max}}$  (i.e., the largest value of  $R(t)$  in  $\ell$  independent realizations). For  $\rho(x) \sim x^{-(1+\mu)}$ , the tails of the distribution of  $R(t)$  behave as the tails of  $\rho(x)$  and therefore, from extreme value theory, we expect the scaling

$$R^{\text{max}} \sim N^{1/\mu}, \quad (10)$$

for  $0 < \mu < 2$  and  $R^{\text{max}} \sim \sqrt{N}$  for  $\mu > 2$  (including the Poisson distribution).

## Data availability

Raw data – i.e., tweet IDs as per Twitter policies – are available from the authors upon request. Anonymized network and activity data used in this study are made publicly available at the: <https://github.com/manlius/SocialBursts>.

Received: 16 October 2019; Accepted: 17 February 2020;

Published online: 13 March 2020

## References

1. Bagrow, J. P., Wang, D. & Barabasi, A. L. Collective response of human populations to large-scale emergencies. *PLoS one* **6**(3), e17680 (2011).
2. Barabasi, A. L. The origin of bursts and heavy tails in human dynamics. *Nature* **435**(7039), 207 (2005).
3. Roehner, B., Sornette, D. & Andersen, J. V. Response functions to critical shocks in social sciences: An empirical and numerical study. *International Journal of Modern Physics C* **15**(06), 809–834 (2004).
4. Lerman, K. & Ghosh, R. Information contagion: An empirical study of the spread of news on digg and twitter social networks in *Fourth International AAAI Conference on Weblogs and Social Media*. (2010).
5. Bakshy, E., Hofman, J. M., Mason, W. A. & Watts, D. J. Everyone's an influencer: quantifying influence on twitter in *Proceedings of the fourth ACM international conference on Web search and data mining*. (ACM), pp. 65–74 (2011).
6. Wu, S., Hofman, J. M., Mason, W. A. & Watts, D. J. Who says what to whom on twitter in *Proceedings of the 20th international conference on World Wide Web*. (ACM), pp. 705–714 (2011).
7. Baños, R. A., Borge-Holthoefer, J. & Moreno, Y. The role of hidden influentials in the diffusion of online information cascades. *EPJ Data Science* **2**(1), 6 (2013).
8. Goel, S., Anderson, A., Hofman, J. & Watts, D. J. The structural virality of online diffusion. *Management Science* **62**(1), 180–196 (2015).
9. Crane, R. & Sornette, D. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences* **105**(41), 15649–15653 (2008).
10. Crane, R., Schweitzer, F. & Sornette, D. Power law signature of media exposure in human response waiting time distributions. *Physical Review E* **81**(5), 056101 (2010).
11. Mathews, P., Mitchell, L., Nguyen, G. & Bean, N. The nature and origin of heavy tails in retweet activity in *Proceedings of the 26th International Conference on World Wide Web Companion*. pp. 14930–1498 (2017).
12. Hawkes, A. G. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* **58**(1), 83–90 (1971).
13. Szabo, G. & Huberman, B. A. Predicting the popularity of online content. *Communications of the ACM* **53**(8), 80–88 (2010).
14. Ratkiewicz, J., Fortunato, S., Flammini, A., Menczer, F. & Vespignani, A. Characterizing and modeling the dynamics of online popularity. *Physical Review Letters* **105**(15), 158701 (2010).
15. Borghol, Y. et al. Characterizing and modelling popularity of user-generated videos. *Performance Evaluation* **68**(11), 1037–1055 (2011).
16. Figueiredo, F., Benevenuto, F. & Almeida, J. M. The tube over time: characterizing popularity growth of youtube videos in *Proceedings of the fourth ACM international conference on Web search and data mining*. (ACM) pp. 745–754 (2011).
17. De Vries, L., Gensler, S. & Leeflang, P. S. Popularity of brand posts on brand fan pages: An investigation of the effects of social media marketing. *Journal of interactive marketing* **26**(2), 83–91 (2012).
18. Bandari, R., Asur, S., Huberman, B. A. The pulse of news in social media: Forecasting popularity in *Sixth International AAAI Conference on Weblogs and Social Media* (2012).



19. Pinto, H., Almeida, J. M. & Gonçalves, M. A. Using early view patterns to predict the popularity of youtube videos in *Proceedings of the sixth ACM international conference on Web search and data mining*. (ACM), pp. 365–374 (2013).
20. Naaman, M., Becker, H. & Gravano, L. Hip and trendy: Characterizing emerging trends on twitter. *Journal of the American Society for Information Science and Technology* **62**(5), 902–918 (2011).
21. Lehmann, J., Gonçalves, B., Ramasco, J. J. & Cattuto, C. Dynamical classes of collective attention in twitter in *Proceedings of the 21st international conference on World Wide Web*. (ACM), pp. 251–260 (2012).
22. Omodei, E., De Domenico, M. & Arenas, A. Characterizing interactions in online social networks during exceptional events. *Frontiers in Physics* **3**, 59 (2015).
23. Dorogovtsev, S. N., Goltsev, A. V. & Mendes, J. F. Critical phenomena in complex networks. *Reviews of Modern Physics* **80**(4), 1275 (2008).
24. Wu, F. & Huberman, B. A. Novelty and collective attention. *Proceedings of the National Academy of Sciences* **104**(45), 17599–17601 (2007).
25. Rizoiu, M. A. *et al.* Expecting to be hip: Hawkes intensity processes for social media popularity in *Proceedings of the 26th International Conference on World Wide Web*. (International World Wide Web Conferences Steering Committee), pp. 735–744 (2017).
26. Gleeson, J. P., O’Sullivan, K. P., Baños, R. A. & Moreno, Y. Effects of network structure, competition and memory time on social spreading phenomena. *Physical Review X* **6**(2), 021019 (2016).
27. Myers, S. A. & Leskovec, J. The bursty dynamics of the twitter information network in *Proceedings of the 23rd international conference on World wide web*. (ACM), pp. 913–924 (2014).
28. De Domenico, M., Lima, A., Mougél, P. & Musolesi, M. The anatomy of a scientific rumor. *Scientific reports* **3**, 2980 (2013).
29. He, X. & Lin, Y. R. Measuring and monitoring collective attention during shocking events. *EPJ Data Science* **6**(1), 30 (2017).
30. Tannenbaum, N. R. & Burak, Y. Theory of nonstationary hawkes processes. *Physical Review E* **96**(6), 062314 (2017).
31. Fujita, K., Medvedev, A., Koyama, S., Lambiotte, R. & Shinomoto, S. Identifying exogenous and endogenous activity in social media. *Physical Review E* **98**(5), 052304 (2018).
32. Miotto, J. M., Kantz, H. & Altmann, E. G. Stochastic dynamics and the predictability of big hits in online videos. *Phys. Rev. E* **95**(3), 032311 (2017).
33. Bouchaud, J. P. & Georges, A. Anomalous diffusion in disordered media: Statistical mechanisms, models and physical applications. *Physics Reports* **195**, 127–293 (1990).
34. Kwak, H., Lee, C., Park, H. & Moon, S. What is twitter, a social network or a news media? in *Proc. 19th Intern. Conf. on World Wide Web*. (ACM), pp. 591–600 (2010).

## Acknowledgements

M.D.D. acknowledges partial financial support from the Max Planck Institute for the Physics of Complex Systems (Visitors program 2016). EGA was funded by the University of Sydney bridging Grant G199768 and CTDS incubator scheme.

## Author contributions

M.D.D. and E.G.A. designed the work, wrote the main manuscript text and have approved the submitted version.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41598-020-61523-z>.

**Correspondence** and requests for materials should be addressed to M.D.D.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020