An Agent-based Model of Posting Behavior During Times of Societal Unrest

Krishna C. Bathina[1] ⋆, Aruna Jammalamadaka[2] Jiejun Xu[2], and Tsai-Ching Lu[2]

[1] Indiana University, Bloomington IN 47408, USA,
bathina@indiana.edu,
[2] HRL Laboratories, LLC, Malibu, CA 90265,
{ajammalamadaka, jxu, tlu}@hrl.com

Abstract. Social media is increasingly monitored during periods of societal unrest to gauge public response and estimate the duration and severity of related protest events. To this end, we build an agent-based simulation model that accurately describes the shift in posting behavior of social media users as related to a real historical event. Specifically, we focus on the indicators and appropriate representation for an agent to become an "activist", or someone who disseminates protest-related posts during times of unrest. We validate our model using a complete collection of Tumblr data from 6 months prior to the Ferguson protest of 2014, until the state of emergency was lifted. We then build an agent-based model based on parameters estimated from before and during the protest. Our model is validated by the similarity of distributions of calculated metrics to the empirically observed data and can be used to predict the posting behavior of protesters. The model is extensible to quantifying behavioral deviations, studying information dynamics, and predicting the effect of real world events on social media behavior.

Keywords: agent based model, information diffusion, social media, political unrest, Tumblr

1 Introduction

"Hashtag activism", first mentioned in reference to the Occupy Wall Street movement, has been increasingly popular during recent protests [2], including viral hashtags such as #OccupyWallStreet, #SOSVenezuela, #HKClassBoycott, and #StopLieAboutTurkey. This is due to both increased awareness of global news events and increased use of social media platforms such as Twitter, Facebook, Instagram, and Tumblr. These factors, along with the low cost of posting, enable activists to easily further protest issues to populations that are not directly affected by the protest [13, 14]. While the efficacy of these protests have come into question [15], there is a positive correlation between social media usage and

⋆ This work was conducted while the first author was doing an internship at HRL Laboratories, LLC.

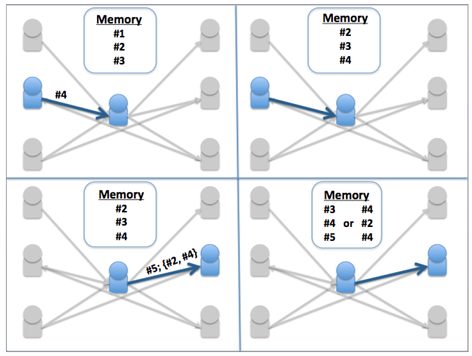

Fig. 1: A cartoon visualization of the agent-based model for online social media users before a protest event.

political participation [16], therefore the spreading behavior of these hashtags is of interest to those studying information dynamics during times of civil unrest.

Agent-based models (ABMs) are computational models with autonomous agents, an environment, and mechanistic behaviors that can be used to represent and simulate emergent behavior from complex, non-linear mathematical systems [6, 4]. Recent advances in computing and access to unprecedented amounts of data have allowed the development of realistic, data-driven agent-based social media models built with many complicated cognitive features. For example, the attention span of an agent is important for human dynamic models, especially with large and multiple flows of information [12]. Another important factor is the natural decrease in propagation of information and its exponential decay with time [18]. In [3], an ABM is built to simulate the dynamics of an insurgent population. Using people as agents and simple mechanistic rules, their model successfully provides an explanation to conflict data across insurgencies. In this paper, we focus on building an ABM that accurately models social media behavior during an actual protest by switching rules for agents during the protest depending on whether or not they become an "activist". A cartoon representation of our model is shown in Figure 1. In the upper left corner, an agent (i.e., an online social media user) chooses to adapt a meme #4 from its neighbor. Subsequently, in the upper right corner, the same agent updates its memory so that the oldest meme is removed and #4 is added. In the lower left corner, the agent can choose to post either a novel meme (e.g.,#5) or chose existing memes from memory (e.g., #2, #4). Finally, in the bottom right quadrant, similar to above, the memory of the agent is updated, either with #5 or with #2 and #4. Such a diffusion style is inspired by the work in [17].

We choose the Ferguson protest of 2014 as the topic of our case study. On August 9th, Michael Brown, an African American resident of Ferguson, Missouri, was shot dead by now ex-Police Officer Darren Wilson [1]. Ferguson locals started protesting later that day and were met with officers in riot gear. The tensions escalated and included a declared state of emergency by the governor, ferguson, mikebrown, michaelbrown, tcot, justiceformikebrown, handsupdontshoot, mediablackout, darrenwilson, dontshoot, fergusonshooting, opferguson, policestate, missouri, ripmikebrown, handsup, police, iftheygunnedmedown, fergusonriot, prayforferguson, stlouis, standwithferguson, blacklivesmatter, nojusticenopeace, peaceinferguson, fergusonpd, whereisjustice, occupyferguson, ccot, arrestdarrenwilson, officergofuckyourself, curfew, crimebutnotime, fergusonpolice, fergusononfireusa, michael, brown, humanrightsferguson, michealbrown, mikebrownfuneral, blackyouthmatter, fergusonsolidarity, militarizationofpolice, justiceformichael, justiceformichaelbrown, mikebrownrally, militarizedpolice, dcferguson, stoppolicebrutality, policemilitarization, pleasedontshoot, fergusonqs, sosferguson, copwatch, resistwithferguson, mikebrownnola, furgeson, iammikebrown, fergusonscanner, protests, endpoliceterror, badgecam, fergusonriots, feedferguson, direnferguson, violenceincites, fergusonlive, freeferguson, endpolicebrutality, fergusontapes, fergusoncoverup, nojustice, fergson, standup, justiceformike

Fig. 2: Memes (Hashtags) related to the Ferguson protest.

a state mandated curfew, intervention by the National Guard and increased violence on both sides. On September 3rd 2014, the state of emergency was lifted. While the physical protests were local, it gained national popularity due to increasing social media coverage. Hashtags such as #Ferguson, #fergusondecision, #blacklivesmatter, and #justiceformikebrown became instantly viral. An estimate from Twitter showed more than 3.5 million tweets mentioning Ferguson were published within 3 hours after the grand jury decision not to indict Wilson [9].

The rest of the paper is outlined as follows: The Methods section describes the data pre-processing and extraction, analysis methods, and the ABM itself. This is followed by the Results section in which we describe how our model compares to the empirically observed data using various network metrics. Finally the paper concludes with possible extensions and future use cases.

2 Data Collection

We used data from Tumblr for our analysis. Tumblr is a Yahoo! owned microblogging, social media website [7]. Users upload blog posts that can include text, images, and video. Users can also unilaterally follow other users and reblog their content. Unlike Twitter, Tumblr does not have a character limit on posts and thus is more commonly viewed as a social media for themed blogs. Our dataset contains every single blog post and re-blog from 2012 to 2014. Each entry contains the original post, the memes (hashtags) used in the post, the date of the re-blog, the original source of the re-blog (root), and the direct source of the re-blog (parent). For our experiment, we first gathered a list of memes that were in support of Ferguson as defined by [11]. The memes are shown in table 2.

We then define two time periods in our dataset and for our model: May to August 8th as *Before* the protest and August 9th to September 3rd as *During* the protest. The *Before* corresponds to three months before the protest for modeling non-protest behavior and the *During* corresponds to the period from the killing of Michael Brown to the day the national emergency was lifted for modeling protest behavior. We then found all users that have used any of the above protest memes at least once in the *During* period. Once the user population was collected, we extracted all of their posts and re-blogs from both the *Before* and *During* periods. We label all memes that are not in the list above as *non-protest memes*. We also extracted all posts and re-blogs in the same time period from 10,000 random users that never used one of the above protest memes as a control group, in order to first test if there is truly a statistical difference in behavior between this group and the protest-meme-using group.

Typical social media APIs only allow for a partial data collection, or implement waiting times that make the collection of a complete dataset very difficult. The public Twitter API allows 180 queries per 15 minute window, Twitter Decahose only collects 10% of the total data, and Instagram allows a maximum collection of 5000 posts per hour. Our case study differs in that we have the full dataset for our chosen period of interest, and thus our model can be accurately validated. In total, this period contains 220 million posts and 764 million memes. During the protest, about 1.7% of the posts and 2.1% of the tags were about the Ferguson protest. From this dataset, we are able to extract the full re-blog network and analyze every blog and re-blog. This network consists of 413,867 nodes and around 23 million total edges.

2.1 Metrics

In order to describe our data, we use four different metrics for quantifying social media behavior, as defined by [17]. Meme-centric posting metrics include the **Meme Time** and the **Meme Popularity**. Meme Time is the longest consecutive number of days that a meme was posted in the dataset and the Meme Popularity is the average number of posts of a meme per day.

Agent-centric posting metrics include the **User Entropy** and the **User Attention**. User Entropy is the average Shannon Entropy of the memes posted by a given user per day, and is given by $H(X) = -\sum_{i=i}^{n} P(x_i)\ln(P(x_i))$ where x_i represents a given meme, and n is the total number of memes posted by that user that day. User Attention is calculated as the average number of re-blogs per user per day. The Meme Popularity, User Entropy, and User Attention are averaged over only days that had posts; days without any posts were ignored.

2.2 Preliminary Analysis

In order to determine whether a change in ABM rules during the protest was needed we performed 4 different preliminary statistical analyses, results of which are shown in Table 1. Protesters are individuals who posted at least one protest meme during the time period of the study. Non-protesters, which are only used for analysis A, were chosen by finding all users that did not use any of the protest hashtags, followed by randomly sampling a set of 10,000 users to prevent any bias. $\Delta\widetilde{x}$ is the difference in median between both groups, or the effect size, and Z is the test statistic from the Kolmogorov-Smirnov test. The *Meme Time* was normalized over the total number of days, allowing the continuous assumptions

	Meme-Centric				User-Centric			
	Popularity		Time		Attention		Entropy	
	$\Delta\widetilde{x}$	Z	$\Delta\widetilde{x}$	Z	$\Delta\widetilde{x}$	Z	$\Delta\widetilde{x}$	Z
A. Non-protesters during vs protesters during	-0.19	0.02	0.07	0.10	-12.79	0.28	-2.14	0.31
B. Protesters before vs protesters during	-0.05	0.03	-0.03	0.89	-10.65	0.01	-1.63	0.02
C. Non-protest memes during vs protest memes during	2.35	0.78	0.52	0.89				
D. Non-protest memes before vs protest memes during	12.95	0.82	0.47	0.89				

Table 1: Preliminary analysis showing statistical differences in posting behavior as captured by the four metrics.

of the KS test to hold. P-values are not reported because $-\log p > 30$ for each test, and thus were significant. For most comparisons, the $\Delta\widetilde{x}$ were very small. Analyses C and D, on the other hand, show that the protest memes had much more *Popularity* than all non-protest memes. Also, Analyses A and B show that the *User Attention* for both non-protesters and protesters before the protest was larger than protesters during the protest. Our results show that a difference exists between all compared groups, especially the *Popularity* of memes and the *Attention* of users.

3 Model Description

Our model is meant to mimic the natural posting patterns and influence of connected users during protest and non-protest periods. The model consists of Tumblr users as agents and the full re-blog network as their environment, where directed edges represent the *flow* of memes. Each time step in our model represents one day. The total number of posts in the "Before" and "During" simulation periods are equal to the observed total number of posts during those periods, and an equal number of posts occur on each day of the simulation. Agents have a finite-sized *Memory* that contains a list of memes with repetitions. The memory is finite because it better models the limited attention that is evident among social media users [10, 5]. If new memes are added to the memory, the oldest meme is removed from the list, representing the discovery that the number of memes to which a user can pay attention is bound, and therefore the injection and survival of new memes comes at the expense of others [17]. The model contains five parameters; P_n is the probability of posting a novel meme, P_r is the probability of posting multiple memes per post before the protest, P_{rn} is the probability of posting multiple non-protest memes during the protest, and P_{rp} is the probability of posting multiple protest memes during the protest. Finally, P_m represents the proportion of protest memes needed in memory to post about the protest. These probabilities are adapted from [17].

The novel aspects of our model come from splitting the model into two time periods; before and during the protest. A flowchart for the *Before* model is shown in Figure 3a. At initialization, the re-blog network for the data is loaded into the model. We use the largest connected component of the original network, containing 412,803 nodes. The agents' memories are then loaded with random hashtags.

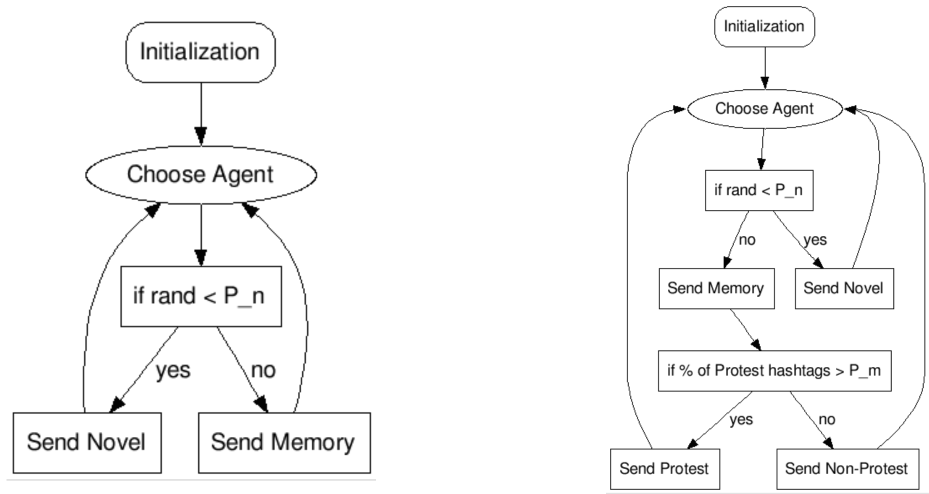

(a) *Before* model. (b) *During* model.

Fig. 3: The above flowcharts represent the algorithms for both the *Before* and *During* models, as described in the text.

At each iteration, an agent is chosen to post with a probability proportional to their out-degree, which has previously been shown to approximate true posting behavior [8]. This agent then either posts a novel hashtag with probability P_n, or posts a set of hashtags from memory. If the agent is posting from memory, each hashtag in memory is added to the post with probability P_r. After every post, the agent's memory, along with the memories of its neighbors are updated with the posted memes as shown in Figure 1.

The *During* model, shown in Figure 3b, is initialized with the agent attributes and network from the end result of the *Before* model. An initial number of agents, equal to the number of actual protestors on the first day of the protest, are randomly chosen as protesters, and protest memes with frequencies proportional to the observed counts on the first day, are added to their memory. The model itself is identical to the *Before* model until an agent chooses to post from memory. If the percentage of protest memes in their memory is greater than P_m, the agent has become an *activist*, and consequently, this agent posts only protest memes. Each protest meme is chosen with probability P_{rp} with each post containing at least one meme. If the percentage in memory is not greater than P_m, the agents posts only non-protest memes, each with probability P_{rn}. To clarify, all agents in our model are protesters; they become activists once more than P_m percent of their memory is filled with protest memes. Again, after every post, the agent's memory along with the memories of its neighbors are updated accordingly.

All macro-level model probabilities were calculated empirically from the data. P_n is calculated by finding the average number of posts with a new meme per

Parameter	Value	Significance
Memory Size	10	Measure of attention span
P_n (Before)	0.2657	Novel meme probability
P_r (Before)	$\frac{2.624}{10}$	Re-blog probability
P_n (During)	0.3097	Novel meme probability
P_m (During)	0.6	Prop. of protest memes in memory to become an "activist"
P_{rp} (During)	$\frac{3.145}{10}$	Protest meme re-blog probability
P_{rn} (During)	$\frac{2.622}{10}$	Non-protest meme re-blog probability

Table 2: Value and significance of each parameter of the agent-based model.

unit time (day). The P_r parameter family is calculated by the average number of memes per post divided by the length of the agent's *Memory*. P_r represents the average number of memes per post before the protest, while P_{rp} represents the average number of memes per post during the protest for posts that include protest memes and similarly, P_{rn} is calculated by the average number of memes per post during the protest for posts that do not include protest memes. P_m and the size of the *Memory* are tunable parameters. The results of all the parameters are shown in Table 2.

4 Results

Overall, our ABM metrics show that results from the *Before* and *During* model are quite similar to empirical results from observed data. In this section we choose to focus on the emergent results from the *During* model because it is the major part of our contribution. These results are shown in the normalized histograms of Figure 4. Figure 4a shows that the *User Attention* from our model did not match that from data as well as we expected. Our model shows a linear distribution of *User Attention* because of our assumption that posting is proportional to the number of out edges. However, the results indicate that the number of Tumblr users with a moderate average number of posts per day are higher than expected. Even with this mismatch, we believe our assumption is reasonable based on previous studies, and we hesitate to overfit our model by incorporating the observed total number of posts per day. The *Entropy* (Figure 4b), of the model did match the data with a slight increase to a peak around 1.0, and then a rapid decrease afterwards, which suggests that most users tended to post with very little variety per day. In the model, *Entropy* is a factor of the rate of novel memes, protest memes, and non-protest memes. Increasing the rate of novel memes, P_n, would increase the average user entropy while increasing the rate of protest and non-protest memes would decrease the entropy.

Figure 4c shows that although the model and data distributions have similar shapes, the model tended to overestimate the *Meme Popularity*. This is most likely due to our posting behavior assumption since *Meme Popularity* is a function of what memes are posted, and thus, which users are posting. But, with such a low difference in probabilities and a similar distribution shape, we believe the deviances are reasonable. The *Meme Time* for the model and data in Figure 4d show similar distributions, with both flattening out as the time increases, suggesting that the majority of Tumblr hashtags are not re-blogged. *Meme Time* is a function of the re-blog parameters, P_r, P_{rp}, P_{rn}; increasing their values would cause an increase in the lifetime of the meme.

The results from the *During* model in Figure 4 may look very similar to the data simply due to a large proportion of non-protest memes. Therefore, to capture the true effect of the model, we show unnormalized histograms for only the protest memes in Figure 5. Figure 5a shows a very similar behavior between *Meme Times* in the "During" model and the data, with the model tending to slightly underestimate the times. Similarly, the *Meme Popularity* in Figure 5b shows that the shape of the model results and data distributions match well, but the model tends to overestimate the popularity by about a factor of 10. Overall, the difference in model and empirical results are small, therefore we believe that our model successfully and accurately describes the full Tumblr dataset.

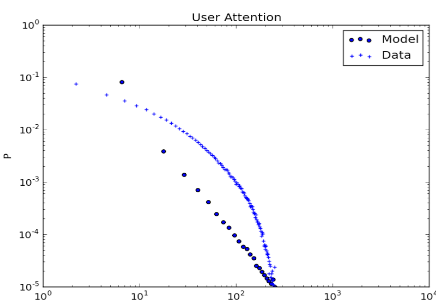

(a) User Attention

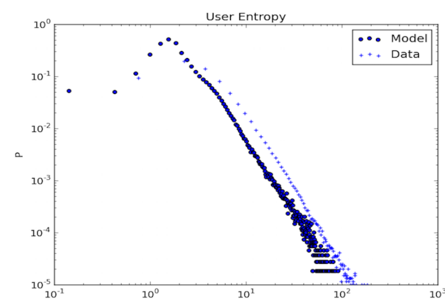

(b) Entropy

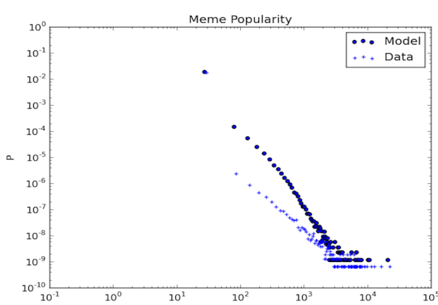

(c) Meme Popularity

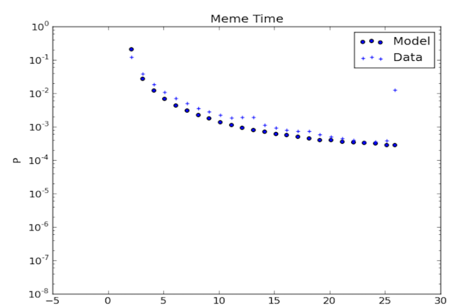

(d) Meme Time

Fig. 4: The above plots show the comparison between model results and observed data via normalized histograms of the defined metrics. All plots are shown on a log-log scale except for (d) Meme Time, which is shown on a linear-log scale.

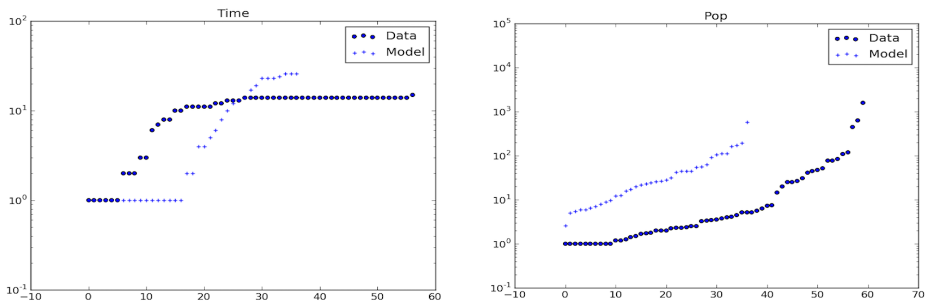

(a) Meme Time

(b) Meme Popularity

Fig. 5: The above plots compare metrics on only the protest memes, using both the *During* model and the observed data. The plots are on a linear-log scale and use unnormalized histograms to highlight the magnitude of protest memes.

5 Discussion

We built a new ABM which uses a real protest event to represent a radical change in behavior of a sub-population of the agents. We then validated our model empirically by analyzing Tumblr data during the Ferguson protest of 2014. We acknowledge that this is an empirical study, and validation on an entirely new protest dataset is required in order for the model to be proven usable for prediction and simulation of future events. We chose not to perform cross-validation by sub-sampling the network used in this study, opting instead for the more realistic analysis of posting interactions that ensues from analyzing the full Tumblr re-blog network. However, we believe that our model, and the extensions of it described below, can still be useful in quantifying and simulation social media posting behavior during times of protest.

There are many straightforward extensions of this model that would enable its use in more complex scenarios. For example, we may want to model the compound, possibly cascading, effects of several consecutive real-world protest events, such as the ongoing protests in Venezuela. An alternate extension involves allowing more than just two types of agents (i.e. activist or non-activist). For example, a class of neutral agents could represent non-protesters and be added to the network. These non-protesting agents interact with the protesting agents and influence the spread of non-protest memes. This flexibility, specifically the ease of implementing heterogeneous agents with diverse rule sets is one of the many strengths of ABMs. However, the more complex the model becomes, the more computationally costly it is as well, therefore this added realism may require a trade-off such as using a randomly sampled subset of the entire network.

Another natural extension is the use of multiple sources of information. The model can include multiple networks of different online social media outlets and will more realistically predict the flow of memes. The problem with these kinds of models is that data across multiple sources is difficult to gather, thus validating the model is also difficult. By validating our model on the full Tumblr dataset, we were able to measure how accurately posting behaviors are captured. We believe this agent-based model can be used by researchers to study information diffusion and the feedback effects between physical events and online social media networks.

References

1. Ferguson unrest (2016), https://en.wikipedia.org/wiki/Ferguson_unrest
2. Augenbraun, E.: Occupy wall street and the limits of spontaneous street protest (2011), https://www.theguardian.com/commentisfree/cifamerica/2011/sep/29/occupy-wall-street-protest
3. Bohorquez, J.C., Gourley, S., Dixon, A.R., Spagat, M., Johnson, N.F.: Common ecology quantifies human insurgency. Nature 462(7275), 911–914 (2009)
4. Bonabeau, E.: Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences 99(suppl 3), 7280–7287 (2002)
5. Boyd, D.: Streams of content, limited attention: The flow of information through social media. Educause Review 45(5), 26 (2010)
6. Bruch, E., Atwell, J.: Agent-based models in empirical social research. Sociological methods & research p. 0049124113506405 (2013)
7. Chang, Y., Tang, L., Inagaki, Y., Liu, Y.: What is tumblr: A statistical overview and comparison. ACM SIGKDD Explorations Newsletter 16(1), 21–29 (2014)
8. Chung, F., Lu, L.: The average distances in random graphs with given expected degrees. Proceedings of the National Academy of Sciences 99(25), 15879–15882 (2002)
9. Doren, J.: Ferguson tweets, hashtags spike as anger rises (2014), http://www.usatoday.com/wlna/news/nation-now/2014/11/26/twitter-ferguson-blacklivesmatter/70132602/
10. Hodas, N.O.: How limited visibility and divided attention constrain social contagion. In: In SocialCom. Citeseer (2012)
11. Jules, B.: Hashtags of ferguson (2014), https://medium.com/on-archivy/hashtags-of-ferguson-8f52a0aced87.vbia14pwu
12. Moussaid, M., Helbing, D., Theraulaz, G.: An individual-based model of collective attention. arXiv preprint arXiv:0909.2757 (2009)
13. Papic, M., Noonan, S.: Social media as a tool for protest. Stratfor Global Intelligence 3 (2011)
14. Shirky, C.: The political power of social media: Technology, the public sphere, and political change. Foreign affairs pp. 28–41 (2011)
15. Taibi, C.: Fox news panel slams bringbackourgirls hashtag activism (2014)
16. Valenzuela, S.: Unpacking the use of social media for protest behavior the roles of information, opinion expression, and activism. American Behavioral Scientist 57(7), 920–942 (2013)
17. Weng, L., Flammini, A., Vespignani, A., Menczer, F.: Competition among memes in a world with limited attention. Scientific reports 2 (2012)
18. Wu, F., Huberman, B.A.: Novelty and collective attention. Proceedings of the National Academy of Sciences 104(45), 17599–17601 (2007)