

Long Trend Dynamics in Social Media

Chunyan Wang¹, Bernardo A. Huberman^{2*}

¹ *Department of Applied Physics, Stanford University, CA, USA* ,

² *Social Computing Lab, HP Labs, Palo Alto, California, USA*

A main characteristic of social media is that its diverse content, copiously generated by both standard outlets and general users, constantly competes for the scarce attention of large audiences. Out of this flood of information some topics manage to get enough attention to become the most popular ones and thus to be prominently displayed as trends. Equally important, some of these trends persist long enough so as to shape part of the social agenda. How this happens is the focus of this paper. By introducing a stochastic dynamical model that takes into account the user's repeated involvement with given topics, we can predict the distribution of trend durations as well as the thresholds in popularity that lead to their emergence within social media. Detailed measurements of datasets from Twitter confirm the validity of the model and its predictions.

I. INTRODUCTION

The past decade has witnessed an explosive growth of social media, creating a competitive environment where topics compete for the attention of users [1, 2]. A main characteristic of social media is that both users and standard media outlets generate content at the same time in the form of news, videos and stories, leading to a flood of information from which it is hard for users to sort out the relevant pieces to concentrate on [3, 4]. User attention is critical for the understand of how problems in culture, decision making and opinion formation evolve [5–7]. Several studies have shown that attention allocated to on-line content is distributed in a highly skewed fashion [8–11]. While most documents receive a negligible amount of attention, a few items become extremely popular and persist as public trends for long a period of time [12–14]. Recent studies have focused on the dynamical growth of attention on different kinds of social media, including Digg [15–17], Youtube [18], Wikipedia [19–21] and Twitter [22]. The time-scale over which content persists as a topic in these media also varies on a scale from hours to years. In the case of news and stories, content spreads on the social network until its novelty decays [15]. In information networks like Wikipedia, where a document remains alive for months and even years, popularity is governed by bursts of sudden events and is explained by the rank shift model [19].

While previous work has successfully addressed the growth and decay of news and topics in general, a remaining problem is why some of the topics stay popular for longer periods of time than others and thus contribute to the social agenda. In this paper, we focus on the dynamics of long trends and their persistence within social media. We first introduce a dynamic model of attention growth and derive the distribution of trend durations for all topics. By analyzing the resonating nature of the content within the community, we provide a threshold

criterion that successfully predicts the long term persistence of social trends. The predictions of the model are then compared with measurements taken from Twitter, which as we show provides a validation of the proposed dynamics.

This paper is structured as follows. In Section 2 we describe our model for attention growth and the persistence of trends. Section 3 describes the data-set and the collection strategies used in the study, whereas Section 4 discusses the measurements made on data-sets from Twitter and compares them with the predictions of the model. Section 5 concludes with a summary of our findings and future directions.

II. MODEL

On-line micro-blogging and social service websites enable users to read and send text-based messages to certain topics of interest. The popularity of these topics is commonly measured by the number of postings about these topics [15, 19]. For instance on Twitter, Digg and Youtube, users post their thoughts on topics of interest in the form of tweets and comments. One special characteristic of social media that has been ignored so far is that users can contribute to the popularity of a topic more than once. We take this into account by denoting first posts on a certain topic from a certain user by the variable First Time Post, (*FTP*). If the same user posts on the topic more than once, we call it a Repeated Post, (*RP*). In what follows, we first look at the growth dynamics of *FTP*.

When a topic first catches people's attention, a few people may further pass it on to others in the community. If we denote the cumulative number of *FTP* mentioning the topic at time t by N_t , the growth of attention can be described by $N_t = (1 + \chi_t)N_{t-1}$, where the χ_t are assumed to be small, positive, independent and identically distributed random variables with mean μ and variance

*Electronic address: bernardo.huberman@hp.com

σ^2 . For small χ_s , the equation can be approximated as:

$$N_t \simeq \prod_{s=1}^t e^{\chi_s} N_0 = e^{\sum_{s=1}^t \chi_s} N_0. \quad (1)$$

Taking logarithms on both sides, we obtain $\log \frac{N_t}{N_0} = \sum_{s=1}^t \chi_s$. Applying the central limit theorem to the sum, it follows that the cumulative count of *FTP* should obey a log-normal distribution.

We now consider the persistence of social trends. We use the variable vitality, $\phi_t = \frac{N_t}{N_{t-1}}$, as a measurement of popularity, and assume that if the vitality of a topic falls below a certain threshold θ_1 , the topic stops trending. Thus

$$\log \phi_t = \log \frac{N_t}{N_{t-1}} = \log \frac{N_t}{N_0} - \log \frac{N_{t-1}}{N_0} \simeq \chi_t. \quad (2)$$

The probability of ceasing to trend at the time interval s is equal to the probability that ϕ_s is lower than a threshold value θ_1 , which can be written as:

$$\begin{aligned} p &= \Pr(\phi_s < \theta_1) = \Pr(\log \phi_s < \log(\theta_1)) \\ &= \Pr(\chi_s < \log(\theta_1)) = F(\log(\theta_1)), \end{aligned} \quad (3)$$

where $F(x)$ is the cumulative distribution function of the random variable χ . We are thus able to determine the threshold value from $\theta_1 = e^{F^{-1}(p)}$ if we know the distribution of the random variable χ . Notice that if χ is independent and identically distributed, it follows that the distribution of trending durations is given by a geometric distribution with $\Pr(L = k) = (1 - p)^k p$. The expected trending duration of a topic, $E(L)$, is therefore given by

$$E(L) = \sum_0^{\infty} (1 - p)^k p \cdot k = \frac{1}{p} - 1 = \frac{1}{F(\log(\theta_1))} - 1. \quad (4)$$

Thus far we have only considered the impact of *FTP* on social trends by treating all topics as identical to each other. To account for the resonance between users and specific topics we now include the *RP* into the dynamics. We define the instantaneous number of *FTP* posted in the time interval t as FTP_t , and the repeated posts, *RP*, in the time interval t as RP_t . Similarly we denote the cumulative number of all posts-including both *FTP* and *RP*-as S_t . The resonance level of fans with a given topic is measured by $\mu_t = \frac{FTP_t + RP_t}{FTP_t}$, and we define the expected value of μ_t , $E(\mu_t)$ as the active-ratio a_q .

We can simplify the dynamics by assuming that μ_t is independent and uniformly distributed on the interval $[1, 2a_q - 1]$. It then follows that the increment of S_t is given by the sum of FTP_t and RP_t . We thus have

$$\begin{aligned} S_t - S_{t-1} &= FTP_t + RP_t = \mu_t FTP_t = \mu_t (N_t - N_{t-1}) \\ &= \mu_t \chi_t N_{t-1}. \end{aligned} \quad (5)$$




FIG. 1: The normal Q-Q plot of $\log(N_{10})$. The straight line shows that the data follows a lognormal distribution with a slightly shorter tail.

And also

$$\begin{aligned} E_\mu(S_t) &= E_\mu(S_{t-1}) + a_q(N_t - N_{t-1}) = \\ &= E_\mu(S_{t-2}) + a_q(N_t - N_{t-2}) + \dots = \\ &= E_\mu(S_0) + a_q(N_t - N_0) = a_q N_t. \end{aligned} \quad (6)$$

We approximate S_{t-1} by $\mu_t N_{t-1}$. Taking back to Eq. 5, we have

$$S_t \simeq \mu_t (\chi_t + 1) N_{t-1} \simeq \mu_t e^{\chi_t} N_{t-1}. \quad (7)$$

From this, it follows that the dynamics of the full attention process is determined by the two independent random variables, μ and χ . Similarly to the derivation of Eq. 3, the topic is assumed to stop trending if the value of either one of the random variables governing the process falls below the thresholds θ_1 and θ_2 , respectively. The probability of ceasing to trend, defined as p^* , is now given by

$$p^* = \Pr(\chi_t < \log(\theta_1)) \Pr(\mu_t < \theta_2) = \frac{\theta_2 - 1}{2(a_q - 1)} p, \quad (8)$$

$p = F(\log(\theta_1))$. The expected value of L_q for any topic q is given by

$$E(L_q) = \frac{2(a_q - 1)}{F(\log \theta_1)(\theta_2 - 1)} - 1. \quad (9)$$

Which states that the persistent duration of trends associated with given topics is expected to scale linearly with the topic users' active-ratio. From this result it follows that one can predict the trend duration for any topic by measuring its user active-ratio after the values of θ_1 and θ_2 are determined from empirical observations.

III. DATA

To test the predictions of our dynamic model, we analyzed data from Twitter, an extremely popular social

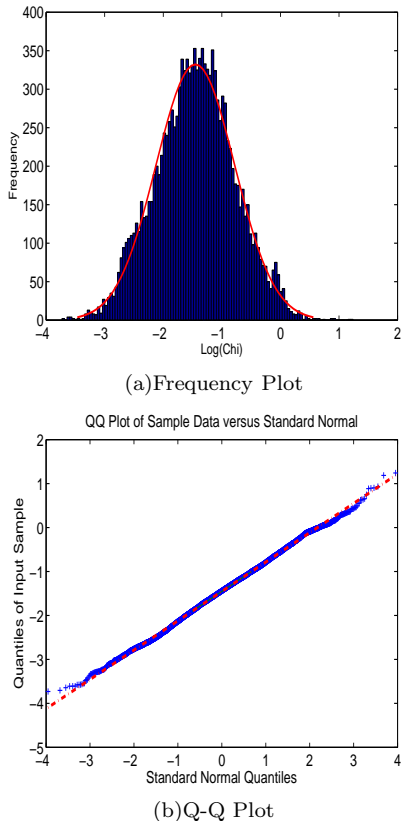


FIG. 2: Distribution of χ over different t and all social trends. $\log(\chi)$ is following a normal distribution with mean equal to -1.4522 and standard deviation equal to 0.6715 . (a) The frequency plot of $\log(\chi)$. (b) The Q-Q plot of $\log(\chi)$

network website used by over 200 million users around the world. Its interface allows users to post short messages, known as tweets, that can be read and retweeted by other Twitter users. Users declare the people they follow, and they get notified when there is a new post from any of these people. A user can also forward the original post of another user to his followers by the re-tweet mechanism.

In our study, the cumulative count of tweets and re-tweets that are related to a certain topic was used as a proxy for the popularity of the topic. On the front page of Twitter there is also a column named trends that presents the few keywords or sentences that are most frequently mentioned in Twitter at a given moment. The list of popular topics in the trends column is updated every few minutes as new topics become popular. We collected the topics in the trends column by performing an API query every 20 minutes. For each of the topics in the trending column, we used the Search API function to collect the full list of tweets and re-tweets related to the topic over the past 20 minutes. We also collected information about the author of the post, identified by a unique user-id, the text of the post and the time of its posting. We thus obtained a dataset of 16.32 million posts on 3361

different topics. The longest trending topic we observed had a length of 14.7 days. We found that of all the posts in our dataset, 17% belonged to the *RP* category.




FIG. 3: The linear scaling relationship between R_n and $\log(t)$ of topic 'Kim Chul Hee', a Korean pop star. The topic kept trending for 14 days on Twitter in September 2010. The number of records that have occurred up to time t scales linearly with $\log(t)$.

IV. RESULTS

From the data-set that we analyzed, we found out that for a fixed time interval of 200 minutes at $t = 10 N_t$ all trends follows a log-normal distribution. As can be seen from Figure 1, the normal Q-Q plot of $\log(N_{10})$ follows a straight line. Different values of t yield similar results. The Kolmogorov-Smirnov normality test of $\log(N_{10})$ with mean 3.5577 and standard deviation 0.3266 yields a P-value of 0.0838. At a significance level of 0.05, the test fails to reject the null hypothesis that $\log(N_{10})$ follows normal distribution, a result which is consistent with Equation 1.

We also observed that the distribution of χ from $\chi_t = \frac{N_t}{N_{t-1}} - 1$. $\log(\chi)$ follows a normal distribution with mean equal to -1.4522 and a standard deviation value of 0.6715 , as shown in Figure 2. The Kolmogorov-Smirnov normality test statistic gives a high p-value of 0.5346. The mean value of χ is 0.0353, which is small for the approximations in Equation 1 and Equation ?? to be valid. We also examined the record breaking values of vitality, $\phi_t = \chi_t + 1$, which signal the behavior of the longest lasting trends. From the theory of records, if the values ϕ_t come from an independent and identical distribution, the number of records that have occurred up to time t , defined as R_n , should scale linearly with $\log(t)$ [24, 25]. As is customary, we say that a new record has been established if the vitality of the trend at the moment is longer than all of the previous observations. As can be seen from Figure 3, this linear scaling is observed over a wide variety of topics. One implication of this observa-

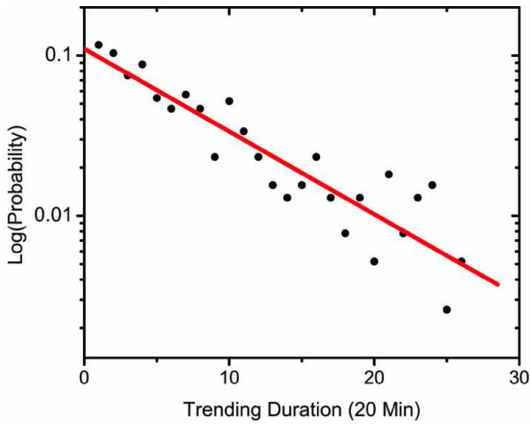


FIG. 4: Semi-log plot of trending duration density. The straight line suggests an exponential family of the trending time distribution. The red line gives a fitting with R-square 0.9112.




FIG. 5: Density plot of trending duration in log-log scale. The distribution of duration deviates from a power law.

tion is that confirms the validity of our assumption that the values of $\chi_1, \chi_2, \dots, \chi_t$ are independent and identically distributed.

Next we turn our attention to the distribution of durations of long trends. As shown in Figure 4 and Figure 5, a linear fit of trend duration as a function of density in a logarithmic scale suggests an exponential family, which is consistent with Eq. 4. The R-square of the fit has a value of 0.9112. From the log-log scale plot in Figure 5, we observe that the distribution deviates from a power law, which is a characteristic of social trends that originate from news on social media [23]. From the distribution of trending times, p is estimated to have a value of 0.12. Together with the measured distribution of χ and Eq. 3, we can estimate the value of θ to be 1.0132.

We can also determine the expected duration of trend times stemming from the impact of active-ratio. The frequency count of active-ratios over different topics is shown in Figure 6, with a peak at $a_q = 1.2$. As can be seen in Figure 7, the trend duration of different topics scales linearly with the active-ratio, which is consistent




FIG. 6: Frequency count of active-ratio over all topics. The maximum ratio is 1.2 among all topics.




FIG. 7: Linear relationship between trending duration and active-ratio, in good agreement with the predictions of model.

with the prediction of Eq. 9. The R-square of the linear fitting has a value of 0.98664. From the slope of the linear fit and $\theta_1 = 1.0132$, and Eq. 9 we obtain a value for $\theta_2 = 1.153$. With the value of θ_1 and θ_2 , we are able to predict the expected trend duration of any given topic based on measurements of its active-ratio.

V. DISCUSSION AND CONCLUSION

In this paper we investigated the persistence dynamics of trends in social media. By introducing a stochastic dynamic model that takes into account the user's repeated involvement with given topics, we are able to predict the distribution of trend durations as well as the thresholds in popularity that lead to the emergence of given topics as trends within social media. The predictions of our

mode were confirmed by a careful analysis of a data from Twitter. Furthermore, a linear relationship between the resonance level of users with given topics, and the trending duration of a topic was derived. The predictive power of this model provides a deeper understanding the popularity of on-line contents. Possible refinements may include the effect of competition between topics, sudden burst of events, the effect of marketing campaigns, or any combination of them. In closing, we note that although the focus in this paper has been on trend dynamics that are featured on social media websites, the framework and model may be suitable to other types

of content and off-line trends. The issue raised - that is, trending phenomenon under the impact of user's repeated involvement - is therefore a general one and should provide ample opportunities for future work.

VI. ACKNOWLEDGMENTS

We acknowledge useful discussions with S. Asur and G. Szabo. C.W. would like to thank HP Labs financial support.

-
- [1] M. E. McCombs and D. L. Shaw (1993) The evolution of agenda setting research: twenty five years in the marketplace of ideas. *Journal of Communication* 43(2):68-84.
 - [2] Josef Falkinger (2008) Limited Attention as a Scarce Resource in Information-Rich Economies. *Economic Journal*, 118(532):1596-1620.
 - [3] Eugene Agichtein, Carlos Castillo, Debora Donato, Aristides Gionis and Gilad Mishne (2008) Finding high-quality content in social media. In Proceedings of the international conference on Web search and web data mining (WSDM), 183-194.
 - [4] Andreas M. Kaplan, Michael Haenlein (2010) Users of the world, unite! The challenges and opportunities of Social Media. *Business Horizons* 53(1):59-68.
 - [5] J. Zhu (1992) Issue competition and attention distraction: A zero-sum theory of agenda setting. *Journalism Quarterly*, 69:825836.
 - [6] Wuchty Stefan, Jones Benjamin and Uzzi Brian (2007) The Increasing Dominance of Teams in Production of Knowledge. *Science*, 316(5827):1036-1039.
 - [7] Guimer Roger, Uzz Brian, Spiro Jarrett and Amaral Lus A. Nunes (2005) Team Assembly Mechanisms Determine Collaboration Network Structure and Team Performance. *Science* 308(5722):697-702.
 - [8] Bernardo A. Huberman, Peter L. T. Pirolli, James E. Pitkow and Rajan M. Lukose (1998) Strong Regularities in World Wide Web Surfing. *Science* 280 (5360): 95-97.
 - [9] Johansen A and Sornette D (2000) Download relaxation dynamics on the WWW following newspaper publication of URL. *Physica A*, 276(1-2):338-345.
 - [10] Bernardo A. Huberman (2001) *The Laws of the Web: Patterns in the Ecology of Information*. The MIT Press, Massachusetts.
 - [11] Vazquez A, et al. (2006) Modeling bursts and heavy tails in human dynamics. *Phys. Rev. E*, 73:036127.
 - [12] W. Russell Neuman (1990) The threshold of public attention. *Public Opinion Quarterly*, 54:159-176.
 - [13] Arjo Klamer and Hendrik P. Van Dalen (2002) Attention and the art of scientific publishing. *Journal of Economic Methodology*, 9(3):289-315.
 - [14] Sitaram Asur, Bernardo A. Huberman, Gabor Szabo, Chunyan Wang (2011) Trends in Social Media: Persistence and Decay. Proceedings of 15th International Conference on Weblogs and Social Media (ICWSM).
 - [15] Fang Wu and B.A. Huberman (2007) Novelty and collective attention. *Proc. Natl. Acad. Sci.* 105:17599.
 - [16] Kristina Lerman and Tad Hogg (2010) Using a Model of Social Dynamics to Predict Popularity of News. Proceedings of 19th International World Wide Web Conference (WWW).
 - [17] J. Leskovec, L. Backstrom and J. Kleinberg (2009) Meme-tracking and the Dynamics of the News Cycle. International Conference on Knowledge Discovery and Data Mining (KDD).
 - [18] R. Crane and D. Sornette (2008) *Proc. Natl. Acad. Sci.* 105:15649.
 - [19] Jacob Ratkiewicz, Santo Fortunato, Alessandro Flammini, Filippo Menczer and Alessandro Vespignani (2010) Characterizing and Modeling the Dynamics of Online Popularity. *Physical Review Letters*, 105:158701.
 - [20] A. Capocci et al. (2006) Preferential attachment in the growth of social networks: The internet encyclopedia Wikipedia. *Phys. Rev. E*, 74:036116.
 - [21] V. Zlatic et al. (2006) Wikipedias: Collaborative web-based encyclopedias as complex networks. *Phys. Rev. E*, 74:016115.
 - [22] B. J. Jansen, M. Zhang, K. Sobel, and A. Chowdhury (2009) Twitter power: Tweets as electronic word of mouth. *J. Am. Soc. Inf. Sci.* 60(11):2169-2188.
 - [23] Sitaram Asur, Bernardo A. Huberman, Gabor Szabo, Chunyan Wang (2011): Trends in Social Media : Persistence and Decay, CoRR, 1102.1402.
 - [24] S. Redner, M. R. Petersen (2006) Role of global warming on the statistics of record-breaking temperatures, *Physical Review E*, 74:061114.
 - [25] Joachim Krug (2007) Records in a changing world, *J. Stat. Mech.*, 07001.

