OPEN

The Simple Rules of Social Contagion

Nathan O. Hodas & Kristina Lerman

SUBJECT AREAS:

COMPUTATIONAL
SCIENCE

USC Information Sciences Institute, Marina del Rey, CA 90292.

INFORMATION TECHNOLOGY

Received
25 September 2013Accepted
24 February 2014Published
11 March 2014

Correspondence and
requests for materials
should be addressed to
N.O.H. (nhodas@isi.
edu) or K.L. (lerman@
isi.edu)

It is commonly believed that information spreads between individuals like a pathogen, with each exposure by an informed friend potentially resulting in a naive individual becoming infected. However, empirical studies of social media suggest that individual response to repeated exposure to information is far more complex. As a proxy for intervention experiments, we compare user responses to multiple exposures on two different social media sites, Twitter and Digg. We show that the position of exposing messages on the user-interface strongly affects social contagion. Accounting for this visibility significantly simplifies the dynamics of social contagion. The likelihood an individual will spread information increases monotonically with exposure, while explicit feedback about how many friends have previously spread it increases the likelihood of a response. We provide a framework for unifying information visibility, divided attention, and explicit social feedback to predict the temporal dynamics of user behavior.

Social media has revolutionized how people create and consume information. Unlike the broadcasts of traditional media, which are passively consumed, social media depends on users to deliberately propagate the information they receive to their social contacts. This process, called social contagion, can amplify the spread of information in a social network. Understanding the mechanics of social contagion is crucial to many applications: creating viral marketing campaigns, evaluating the quality of information, and predicting how far it will spread. While the spread of information is often likened to an infectious disease^{1–4}, social contagion differs in that social media users actively seek out information and consciously decide to propagate it. Because of the constraints of available time and cognitive resources, the ease of discovery will significantly affect information's propensity to go viral^{5,6}. The enormous flux of available social media content often saturates user's ability to process information. In most studies of information propagation on networks, users are considered exposed if they received a message, regardless of whether they see it or not, which can lead to counterintuitive results suggesting that additional exposures inhibit response^{7,8}. In reality, of a user seeing a message depends on how the website arranges content, the flux of incoming information, and the effort the user is willing to expend in discovering information. By accounting for these factors, we demonstrate that social contagion is quite simple and people's responses can be accurately predicted.

From a theoretical perspective, one of the simplest and most widely studied models of social contagion is the independent cascade model (ICM)^{1–4,9,10}. The ICM-class of models assume that each exposure of a healthy (naive) person by an infected (informed) friend leads to an independent chance of information transmission. Therefore, the probability that a healthy individual becomes infected increases monotonically with the number of exposures, potentially causing a global epidemic involving a substantial fraction of the population^{11,12}. However, studies of information spread in social media have identified social behaviors that qualitatively differ from predictions of the ICM. For example, when measuring how people respond to their friends' use of certain memes or recommendations for news articles, repeated exposure initially increases infection probability, but eventually exposure appears to be inhibitory^{7,8}, violating the central assumptions of the ICM. A number of explanations have been offered for this aberration, including complex contagion^{13–15}. In complex contagion, the probability to adopt a behavior, or an idea, varies with the extent of exposure, suggesting that social phenomena may drive response and interact non-trivially with network structure^{16–18}. An alternative explanation invokes the linear threshold model, in which the proportion of friends (past a certain threshold) adopting a behavior determines contagion^{2,19,20}. Among other factors thought to affect social contagion are the novelty²¹ or persistence⁷ of information, and competition with other information⁶. The role of cognitive constraints in online social interactions has not been widely examined, although one study of Twitter demonstrated that people limit themselves to approximately 150 conversation partners²², a number similar to the bound on human social group size²³.

To compare how visibility and social factors contribute to contagion, we collected data from two online social networks: Digg and Twitter. The microblogging service Twitter allows registered users to broadcast short messages, called tweets, to their followers. A message may contain a URL to external web content. In addition to posting a new message, a user can also retweet an existing message, analogous to forwarding an email. Twitter

users create social links by following other users. Each link is directed: we refer to the followed user as the friend, and the following user as the follower. Upon visiting Twitter, a user is presented with a list containing tweets made by friends, with the most recent tweet (or retweet) at the top.

Social news aggregator Digg leverages opinions of its users to help people discover interesting news stories. Users submit URLs to news stories and vote for, or *digg*, stories submitted by others. Users can follow the activity of others. The social user-interface on Digg shows a user a stream of stories his or her friends recently submitted or voted for. The stream is ordered chronologically by time of earliest recommendation (submission or vote) by a friend, with the most recent newly-recommended story at the top. When a user votes for a story, the recommendation is broadcast to a user's followers. However, additional recommendations do not change the story's relative position in the user's default social stream. Instead, a badge appears next to the story showing how many friends have recommended it. When the story receives enough votes, Digg promotes it to its front page. However, before promotion, it can be found through friends' recommendations or on the newly submitted stories list, which at the time of data collection was receiving tens of thousands of new submissions daily.

We use techniques, originating from non-equilibrium statistical physics, to analyze user behavior on these sites. Our approach enables us to separate the factors of social contagion that are attributable to the visibility of information (i.e., how easily it can be discovered in the user interface of each site) from the factors attributable to social influence. After accounting for these factors, social contagion becomes quite simple: each exposure increases the likelihood of a response, and social signals about the number of friends who have previously adopted the information (when such signals are provided by the web site) further amplify response. We demonstrate that we are able to accurately forecast an individual's behavior in real-time on both sites.

Results

Using URLs as markers, we study the spread of information through the follower graphs of Digg and Twitter. A user may be exposed multiple times by friends to a URL. The exposure response function gives the probability of an infection as a function of the number of such exposures. An exposure is defined to occur when a message containing the URL arrives in the user's stream, even if the user does not consciously see it. When aggregated over all users, both Twitter and Digg exposure response functions suggest complex contagion⁷: while initial exposures increase infection probability, further exposures appear to saturate (Twitter) or suppress (Digg) further infection (Fig. 1a). Aggregated exposure response obscures heterogeneous behavior, because it conflates the response of users with different

cognitive loads, i.e., different quantities of information in their stream. A large volume of incoming information, which scales with the number of friends a user follows as $n_f^{1,14}$, reduces the user's ability to find any specific message^{24,25}. The likelihood a user will find a message containing the URL depends on n_f , denoted $\mathcal{P}(n_f)^5$. However, disaggregating only partially ameliorates complications due to underlying heterogeneity; although plotting infection as a function of the fraction of friends adopting the URL on Twitter displays remarkable consistence between user groups (Fig. 2 in⁵), a similar plot using Digg data (Fig. 1b) suggests the contradictory and confusing result that even small increases in exposure dramatically suppress infection. Although a linear threshold model may be consistent with Fig. 1a, neither the ICM nor linear threshold model can simultaneously account for observed trends on Twitter and Digg.

To resolve this contradiction, consider the process of infection on each site. To become infected, a user must first discover at least one message containing the URL. The likelihood the user will see a specific message depends on its position in the user's stream. We use 'visibility' to refer to this quantity. A new message starts at the top of the queue, where it is more likely to be seen because users usually start browsing from the top of a page²⁶. With time, newer messages push it down the queue, where a user is less likely to see it^{27,28}. We approximate a message's dynamic visibility using the time response function, $\mathcal{T}(\Delta t, n_f)$, the probability that a user with n_f friends retweets or votes at a time Δt after the exposure⁵. We plot $\mathcal{T}(\Delta t, n_f)$ for Digg and Twitter in Fig. 2a and 2b, respectively, demonstrating that the visibility of a new message decays rapidly in time. Digg stories were only followed until promotion, which occurs at most 24 hours after appearing on Digg. The data are smoothed using progressively wider smoothing windows, as in⁵.

A model describing user response to multiple exposures must consider the visibility of each exposure. In addition, a website's use of any social signals — for example, displaying the number of friends who recommended the URL — may alter user response, given that they have found the URL. The probability that a user with n_f friends will be infected after n_e exposures is

$$P(t; n_e, n_f) = \sum_{n=1}^{n_e} F(n) V_n(t, \{t_1, \dots, t_{n_e}\}; n_f), \quad (1)$$

where $V_n()$ is the probability of finding n of the n_e exposures occurring at the times $t_1, \dots, t_{n_e}$, and $F(n)$ is the social enhancement factor accounting for the user observing that n of their friends have recommended the story. Note that this formalism averages out content-specific factors and variable weights that a user may ascribe to different friends.

The particular functional form of V_n depends on details of the website user-interface. On Twitter, all messages start at the top of

Figure 1 | The exposure response functions for Twitter and Digg, (a) as a function of total number of votes for the URL a user receives in their information stream, and (b) as a function of the fraction of friends adopting story, for Digg only. The equivalent results for Twitter can be found in⁵.

Figure 2 | The time response functions for (a) Digg and (b) Twitter for different user classes.

the stream. By scanning the stream, a user can discover each message independently, so any of the exposures can result in an infection. This behavior is well approximated by the probability of becoming infected by at least one exposure (see Supplement), given by

$$P_{Tw}(t; n_f, n_e) = P_0 F(n_e) \left[1 - \prod_{i=1}^{n_e} (1 - \mathcal{P}(n_f) \mathcal{T}(\Delta t_i, n_f)) \right] + v_{min}, \quad (2)$$

where v_{min} is the effective minimum visibility of a message in the Twitter interface, the proportionality P_0 is fitted by minimizing weighted mean absolute percent (WMAPE), as described in the Methods, and n_e is the number of exposures to the URL at time t . Underlying activity rates and cultural norms vary from site to site, so P_0 can be interpreted as a task-specific scale factor. The constant v_{min} is due to the ability to discover the URL outside the social media site or via other interfaces.

We calculate $\mathcal{P}(n_f)$ by measuring the average probability of retweeting the URL for users who were exposed only once. The average is taken over all users with n_f friends, as described in^{5,25}. The time response function $\mathcal{T}(\Delta t, n_f)$ describes the visibility after Δt seconds since exposure. Specifically, it is the probability that a user with n_f friends retweets/votes at the indicated interval Δt after a URL's arrival, given that the user votes on that URL.

The Digg user-interface differs from Twitter in that messages are by default ordered by the time of their *first* appearance in the user's stream. Additional votes do not alter its position but are reflected in a badge next to the URL showing the number of friends, n_e , who voted for the URL. The badge provides a social signal, which may alter user response. Because of the user-interface, Eq. (1) reduces to

$$P_{Digg}(t; n_f, n_e) = F'(n_e) (P'_0 \mathcal{P}'(n_f) \mathcal{T}'(\Delta t, n_f) + v'_{min}), \quad (3)$$

where Δt is the time elapsed from the first vote by a friend, and the primes indicate Digg specific values for each quantity. We empirically determined $F'(n_e)$ using a maximum likelihood estimate, described in the Methods. Social feedback in Digg results in large amplification of the probability of infection, shown in Fig. 3c. This could have multiple origins, including endorsement by friends²⁹, or from the increased visibility of the URL via alternative ways of discovering it on Digg, such as sorting URLs by popularity.

To validate the proposed model of social contagion, we forecast user activity and compare it to observed activity. Specifically, we calculate the observed frequency that a user with n_f friends retweeted a URL in our Twitter dataset or voted for one in the Digg dataset in the subsequent 30 seconds. Then, using Eq. (2) or Eq. (3), we calculate the theoretical probability that a user with that many friends would act in those 30 seconds, given the same exposures. Data were divided into a test set and training set. Parameters were estimated on the training set. Results are shown

from the test set. Plotting the predicted versus observed probabilities allows us to graphically assess the accuracy of the contagion model. Unbiased forecasts lie along the unit-slope line. The forecasted responses on Twitter (Fig. 4b) and Digg (Fig. 4d) have a WMAPE error of 0.5% and 1.5%, respectively. Ignoring social enhancement, and thereby utilizing ICM, produces systematically biased results, shown in Figs. 4a and 4c. Without this social enhancement, Twitter and Digg have WMAPE error of 0.7% and 12.2%, respectively. Although we do not know the specific cause of this difference, we may surmise that it is due primarily to the explicit social feedback present on Digg but absent on Twitter. It appears users on Twitter adopted content based primarily on ease of discovery (visibility). Additionally, a model not incorporating visibility decay could not account for variations in user-interface, i.e., Eqs. (2) and (3).

The unbiased fidelity of the proposed model suggests that once visibility of the exposures is taken into account, social contagion operates as a simple contagion, i.e., with infection probability increasing monotonically with the number of exposures. Complex contagion, where “network effects,” appear to play a significant role in the contagion process, may to a large extent be due to the combined factors of visibility and direct social enhancement factors. Moreover, by comparing two different websites with very different user-interfaces, we have demonstrated that it is possible to isolate the factors in social contagion due to social feedback and the user-interface, without directly manipulating the underlying social network or user-interface^{29,30}.

Rapid visibility decay, combined with decreased susceptibility of highly connected users, explains why information in social media fails to spread as widely as predicted by the generic ICM⁸. Although different types of information may spread according to slightly different patterns^{31,32} our analysis is content agnostic, so the reported results are the population average. Explicit social feedback can significantly magnify user response, albeit making it less useful for popularizing high-quality content³³. Unlike Digg, the Twitter user-interface offered no explicit social feedback (beyond trending topics). Users may remember seeing a friend's recommendation of the URL, a factor that could explain the slight social enhancement seen in Twitter response in Fig. 3a. When explicit social feedback is available, as in Digg, Fig. 3d shows that users appear to weigh their actions based on the fraction of friends endorsing a URL instead of considering the absolute number.

Discussion

We show that there are important and surprising differences between the diffusion of information and a disease stemming from cognitive limitations for processing information. In pathogenic contagion³⁴, people with more incoming contacts are more likely to contract a

(a) Twitter, All Users

(b) Twitter, Sub-populations

(c) Digg All Users

(d) Digg, Sub-populations

Figure 3 | The social enhancement factors for Twitter and Digg. (a,c) Averaged over all users, (b,d) Calculated for sub-populations based on n_f . The decay in the social enhancement factor for Twitter can be attributed to residual spam in the dataset.

disease, but in social contagion such people are less likely to become infected: because the volume of information scales with the number of friends a user follows, highly connected users are less likely to notice a particular piece of information, and they require stronger social signals to act (Fig. 3d), on average, than poorly connected users^{5,25}. These highly connected users dominate the high-exposure portion of the average exposure response function (Fig. 1), giving the false impression that more exposures may be counter-productive. Granted, highly connected users tend to be infected earlier³⁵ and also to have more followers²⁴, increasing their influence once they are infected^{17,35}. Users in a tightly connected core of friends may be repeatedly exposed to information, and the present work demonstrates how the combination of social enhancement and awareness contribute to the observed behavior of users in high k-cores participating in larger cascades¹⁷.

By comparing the dynamics of two different websites, we have demonstrated that it is possible to isolate factors in social contagion due to social feedback and the user-interface, without manipulating the underlying social network or user-interface^{16,29,30}. Moreover, the

unbiased fidelity of our model suggests that once visibility of the exposures is taken into account, social contagion operates as a simple contagion, i.e., with infection probability increasing monotonically with the number of exposures. Although our forecasts are only for action within the subsequent 30 seconds, the present work shows that this near-term likelihood can vary by over a factor of 10,000. Although longer periods could be forecast, intervening events, such as receiving additional messages, would invalidate the initial conditions of the forecast. This could be corrected by utilizing higher-order models accounting for the probability of additional messages being received during the forecast window.

Our work highlights how cognitive constraints impact digital content sharing activities. Humans have developed large brains, partly to handle the mental demands of social life^{36,37}, but constraints imposed by our brain's finite information processing bandwidth affect social behavior, for example, by limiting maximum group size²³. Our present results suggest cognitive constraints also affect how individuals utilize information in their dynamic social media streams. Attentive acts, such as browsing a website and reading tweets, require energy;

(a) Twitter, No Social Enhancement

(b) Twitter, Full Model

(c) Digg, No Social Enhancement

(d) Digg, Full Model

Figure 4 | Forecasting accuracy for Twitter and Digg.

because the brain's capacity for mental effort is limited by its energy requirements, so is our attention³⁸. This will reduce responsiveness under conditions of high information load, making explicit social feedback essential for determining the allocation of cognitive resources. Thus, social contagion will be highly dependent on explicit social feedback and the user-interface.

Implicit in our work is the utilization of Big Data as a microscope: we uncover behavioral mechanisms and even the difference between user interfaces. Regardless of the social synergy desired by the website, information discovery costs appear to be an important factor in determining accuracy of activity forecasts. The site's design choice regarding its visibility policy will largely determine the quality of the user experience regarding information discovery and spread. Digg does not refresh the position of information after each recommendation, and the social signals it uses do not compensate for the loss of visibility it suffers over time. Although the current work provides techniques for real-time forecasting of the average user behavior on a specific website, understanding the emergence of globe-spanning viral content will require accounting for the interaction of the

dynamic visibility and social synergy across a multitude of websites and media outlets.

Methods

Data sources. We used Twitter's Gardenhose API to collect tweets over three weeks in Fall of 2010. We retained tweets containing a URL in the message body. We used Twitter's search API to retrieve all tweets containing those URLs, ensuring the complete tweeting history of all URLs, resulting in 3 million tweets in total. We also collected the friend and follower information for all tweeting users, resulting in a social graph with almost 700 K nodes and over 36 M edges. We removed URLs whose retweeting behavior exhibits patterns associated with spam or automatic activity³⁹, leaving us a data set containing 2 K distinct URL's retweeted a total of 213 K times. We use time stamps in tweet metadata combined with the follower graph to track when users are exposed to URLs by a friend and when they retweet them. We define a retweet to be anytime a user tweets a URL that had previously appeared in her Twitter feed. We did not resolve link-shorteners, so different URLs might map to identical content, but we considered each URL to be a unique marker of information. After removing spam URLs, we only consider events where users received a particular URL less than 30 times, to further eliminate likely spam URLs.

We used the Digg API to collect data about 3.5 K stories promoted to the front page in June 2009 and the times at which 140 K distinct users voted for these stories. We also collected information about voters' friends, giving us a social graph with 280 K

users and 1.7 M links. For the present analysis, unless noted otherwise we consider only the voting dynamics occurring before promotion to the front page, so the primary means of information propagation is through the friends interface. Both datasets were divided into training and test sets to rule-out over-fitting in determining the correct interpretation of the data.

To calculate probabilities of response to multiple exposures, the data was broken down into separate time series, each corresponding to the arrival of specific URL-containing tweets or votes into a single user's stream. For each series, at every one-second interval we calculate the quantity we define as 'visibility' of the URL:

$$V_{all}(t, n_f) = 1 - \prod_d (1 - T(t - t_i, n_f)), \quad \text{or} \quad (4)$$

$$V_{first}(t, n_f) = T(t - t_0, n_f), \quad (5)$$

where n_f is the number of friends of the user, $T(\Delta t, n_f)$ is the time response function for a user with n_f friends. V_{all} is proportional to the probability of finding any one of the received messages at time t , while V_{first} is proportional to the probability of finding only the first message.

Data analysis. We calculate $\mathcal{P}(n_f)$ by measuring the average probability of retweeting the URL for users who were exposed once and only once to it. The average is taken over all users with n_f friends, as described in^{5,25}. The time response function $T(\Delta t, n_f)$ describes the visibility of a message since exposure at t_r . This is given by probability, shown in Fig. 2, that a user with n_f friends will retweet a time Δt_i after the exposure, given that retweeting occurred.

The time response function, $T(\Delta t, n_f)$ is produced by calculating, using the observed data, the probability that a user retweets/votes at the indicated interval Δt after a URL's arrival, given that the user votes on that URL. For Twitter data we calculate the time response function only for those events in which a user received the URL once and only once. For Digg, this constraint is lifted, because there are too few such events in the Digg data. The precise time response function depends on n_f because users with many friends receive new messages at a higher rate, causing the visibility of any specific message to decay more quickly⁵. We lack sufficient data to precisely calculate the time response function for each n_f . Instead, we calculated the time response function for users with $n_f = 1-2$, $n_f = 9-11$, and $n_f = 90-110$, producing $\mathcal{T}_1$, $\mathcal{T}_{10}$, and $\mathcal{T}_{100}$, respectively, following the procedure in⁵. To estimate the time response function for arbitrary n_f we interpolated as follows:

$$\begin{aligned} w_1 &\equiv ((n_f - 1)^2 + 10^{-6})^{-1} \\ w_{10} &\equiv ((n_f - 10)^2 + 10^{-6})^{-1} \\ w_{100} &\equiv ((n_f - 100)^2 + 10^{-6})^{-1} \quad \text{for Twitter, or} \\ w_{100} &\equiv (|n_f - 100| + 10^{-6})^{-1} \quad \text{for Digg} \\ \mathcal{T}(\Delta t, n_f) &= \frac{w_1 \mathcal{T}_1(\Delta t) + w_{10} \mathcal{T}_{10}(\Delta t) + w_{100} \mathcal{T}_{100}(\Delta t)}{w_1 + w_{10} + w_{100}}. \end{aligned} \quad (6)$$

To produce the fits for v_{min} and P_0 , we plot the theoretical probability versus the observed probability for an event observed in the data, i.e. forming a function $O(p)$, where p is the calculated probability. That is, no numerical simulations were used, but event timings and the follower network were taken directly from the observed data. We isolated the events corresponding to a receiving a single message, leading to a subset of predictions denoted $O_1(p)$. We then minimize the weighted mean absolute percent error (WMAPE)⁴⁰,

$$\langle |P_0 O_1(p) + v_{min} - p|/p \rangle \quad (7)$$

by searching over P_0 and v_{min} . For Digg, we have $P_0 = 667$, $\log(v_{min}) = -19$. An analytical form for $\mathcal{P}(n_f)$ was determined by fitting to minimize RMS error of the empirically determined $\mathcal{P}(n_f)$ ⁵, giving Digg's $\mathcal{P}(n_f) = A / ((e^{B n_f} + C)(n_f + D)(n_f + E))$, where $A = 7.6 \cdot 10^{-3}$, $B = -6.2 \cdot 10^{-2}$, $C = 1.7 \cdot 10^{-3}$, $D = 3.7$, $E = 17.8$. For Twitter we have $P_0 = 16.6$ and $\log(v_{min}) = -14$, and we used $\mathcal{P}(n_f) = A n_f^P / (n_f + B)$, where $A = 0.3$, $P = 0.16$, $C = 0.55$. Note the E was chosen by minimizing WMAPE error simultaneously with fitting P_0 and v_{min} on the training data, i.e. E 's purpose is to correct for sparsity in the empirically calculated $\mathcal{P}(n_f)$ for Digg.

To calculate the social enhancement factors, we carry out the MLE for $F(n_e)$ in the following manner. We take as axiomatic the true probability of a response given n_e exposures is $F(n_e)P(v)$, where v parameterizes the underlying visibility. Thus, given $N(v)$ observed events for a specific v , the likelihood, ℓ , of observing $N_r(v)$ responses is determined by the binomial distribution $\ell(v, n_e) = \binom{N(v)}{N_r(v)} (F(n_e)P(v))^{N_r(v)} (1 - F(n_e)P(v))^{N(v) - N_r(v)}$. The total log-likelihood of observing the curve $O_{n_e}(v)$ is thus

$$\begin{aligned} \mathcal{L}(n_e) &= \sum_v \log \binom{N(v)}{N_r(v)} \\ &\quad + N_r(v) (\log F(n_e) + \log P(v)) \\ &\quad + (N(v) - N_r(v)) \log (1 - F(n_e)P(v)). \end{aligned} \quad (8)$$

For each value of n_e , we find the value of $F(n_e)$ that maximizes $\mathcal{L}(n_e)$. First, for $n_e = 1$, we define $F(1) = 1$, so we obtain the MLE for $P(v)$ using

$$\frac{\partial}{\partial P(v)} \mathcal{L}(1) = \frac{N_r(v)}{P(v)} - \frac{N(v) - N_r(v)}{1 - P(v)} = 0, \quad (9)$$

giving $P(v) = N_r(v)/N(v)$. Then, for $n_e > 1$, we are left to find the likelihood maximizing $F(n_e)$ given $P(v)$, leading to

$$\frac{\partial}{\partial F(n_e)} \mathcal{L}(n_e) = \sum_v \frac{N_r(v)}{F(n_e)} - (N(v) - N_r(v)) \frac{P(v)}{1 - F(n_e)P(v)}. \quad (10)$$

Numerically solving for $\frac{\partial}{\partial F(n_e)} \mathcal{L}(n_e) = 0$ provides the MLE for $F(n_e)$.

The minimum possible observed probability is bounded by the number of observed events. In the forecasting predictions, the friend-cohort breakdown in Fig. 4 appears to deviate from the observed probabilities at very high and low predicted probabilities. However, this is due to the minimum probability floor rising beyond the predicted = observed line, because events with high visibility and high social influence or very low visibility are less common.

- Newman, M. E. J. Spread of epidemic disease on networks. *Phys. Rev. E* **66**, 016128 (2002).
- Kempe, D., Kleinberg, J. & Tardos, E. Maximizing the spread of influence through a social network. In *Proc. Int. Conf. on Knowledge Discovery and Data Mining*, KDD'03, 137–146 (ACM Press, New York, NY, USA, 2003).
- Gruhl, D., Liben-Nowell, D., Guha, R. & Tomkins, A. Information diffusion through blogspace. *SIGKDD Explor. Newsl.* **6**, 43–52 (2004).
- Anagnostopoulos, A., Kumar, R. & Mahdian, M. Influence and correlation in social networks. In *Proc. Int. Conf. on Knowledge Discovery and Data Mining*, KDD'08, 7–15 (ACM, New York, NY, USA, 2008).
- Hodas, N. O. & Lerman, K. How visibility and divided attention constrain social contagion. In *Proc. ASE/IEEE Int. Conf. on Social Computing*, 249–257 (IEEE Computer Society, Washington, DC, USA, 2012).
- Weng, L., Flammini, A., Vespignani, A. & Menczer, F. Competition among memes in a world with limited attention. *Scientific Reports* **2** (2012).
- Romero, D., Meeder, B. & Kleinberg, J. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter. In *Proc. Int. Conf. on World Wide Web*, WWW'11, 695–704 (ACM, New York, NY, USA, 2011).
- Ver Steeg, G., Ghosh, R. & Lerman, K. What stops social epidemics? In *Proc. 5th Int. Conf. on Weblogs and Social Media*, ICWSM'11, (ACM, New York, NY, USA, 2011).
- Hethcote, H. W. The Mathematics of Infectious Diseases. *SIAM Review* **42**, 599–653 (2000).
- Goldenberg, J., Libai, B. & Muller, E. Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Marketing Letters* 211–223 (2001).
- Castellano, C., Fortunato, S. & Loreto, V. Statistical physics of social dynamics. *Rev. Modern Physics* **81**, 591–646 (2009).
- Satorras, R. P. & Vespignani, A. Epidemic Spreading in Scale-Free Networks. *Phys. Rev. Letters* **86**, 3200–3203 (2001).
- Granovetter, M. The Strength of Weak Ties: A Network Theory Revisited. *Social Theory* **1**, 201–233 (1983).
- Watts, D. J. A simple model of global cascades on random networks. *Proc. Nat. Acad. Sci.* **99**, 5766–5771 (2002).
- Centola, D., Eguiluz, V. & Macy, M. Cascade dynamics of complex propagation. *Physica A* **374**, 449–456 (2007).
- Centola, D. The spread of behavior in an online social network experiment. *Science* **329**, 1194 (2010).
- González-Bailón, S., Borge-Holthoefer, J., Rivero, A. & Moreno, Y. The Dynamics of Protest Recruitment through an Online Network. *Scientific Reports* **1** (2011).
- Rombach, M. P., Porter, M. A., Fowler, J. H. & Mucha, P. J. Core-Periphery Structure in Networks. *arXiv.org* (2012). Retrieved November 2013.
- Granovetter, M. Threshold models of collective behavior. *American J. Sociology* 1420–1443 (1978).
- Granovetter, M. & Soong, R. Threshold Models of Diffusion and Collective Behavior. *J. Mathematical Sociology* **9**, 165–179 (1983).
- Wu, F. & Huberman, B. A. Novelty and collective attention. *Proc. Nat. Acad. Sci.* **104**, 17599–17601 (2007).
- Goncalves, B., Perra, N. & Vespignani, A. Modeling Users' Activity on Twitter Networks: Validation of Dunbar's Number. *PLoS ONE* **6**, e22656 (2011).
- Dunbar, R. I. M. Neocortex size as a constraint on group size in primates. *J. Human Evolution* **22**, 469–493 (1992).

24. Hodas, N. O., Kooti, F. & Lerman, K. Friendship Paradox Redux: Your Friends Are More Interesting Than You. In *Proc. 7th Int. Conf. on Weblogs and Social Media*, ICWSM'13, 1–8 (ACM, New York, NY, USA, 2013).
25. Hodas, N. O. & Lerman, K. Attention and Visibility in an Information-Rich World. In *Int. ICME Workshop on Social Multimedia Research*, 1–6 (IEEE Computer Society, Washington, DC, USA, 2013).
26. Buscher, G., Cutrell, E. & Morris, M. R. What do you see when you're surfing?: using eye tracking to predict salient regions of web pages. In *Proc. 27th Int. Conf. on Human Factors in Computing Systems*, CHI '09, 21–30 (ACM, New York, NY, USA, 2009).
27. Malmgren, R. D., Stouffer, D. B., Campanharo, A. S. L. O. & Amaral, L. A. On Universality in Human Correspondence Activity. *Science* **325**, 1696–1700 (2009).
28. Huberman, B. A., Pirolli, P. L. T., Pitkow, J. E. & Lukose, R. M. Strong Regularities in World Wide Web Surfing. *Science* **280**, 95–97 (1998).
29. Bond, R. M. *et al.* A 61-million-person experiment in social influence and political mobilization. *Nature* **489**, 295–298 (2012).
30. Bakshy, E., Rosenn, I., Marlow, C. & Adamic, L. The Role of Social Networks in Information Diffusion. In *Proc. 21st Int. Conf. on World Wide Web*, WWW'12, 1–10 (ACM, New York, NY, USA, 2012).
31. Yang, J. & Leskovec, J. Patterns of temporal variation in online media. In *Proc. Int. Conf. on Web Search and Data Mining*, WSDM'11, 177–186 (ACM, New York, NY, USA, 2011).
32. Wu, S., Tan, C., Kleinberg, J. & Macy, M. Does Bad News Go Away Faster? In *Proc. Int. Conf. on Weblogs and Social*, ICWSM'11, 646–649 (ACM, New York, NY, USA, 2011).
33. Salganik, M. J., Dodds, P. S. & Watts, D. J. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* **311**, 854–856 (2006).
34. Lloyd-Smith, J. O., Schreiber, S. J., Kopp, P. E. & Getz, W. M. Superspreading and the effect of individual variation on disease emergence. *Nature* **438**, 355–359 (2005).
35. Goldenberg, J., Han, S., Lehmann, D. R. & Hong, J. W. The Role of Hubs in the Adoption Process. *J. Marketing* **73**, 1–13 (2009).
36. Dunbar, R. Evolution of the Social Brain. *Science* **302**, 1160–1161 (2003).
37. Silk, J. B. Social Components of Fitness in Primate Groups. *Science* **317**, 1347–1351 (2007).
38. Kahneman, D. *Attention and effort* (Prentice Hall, 1973).
39. Ghosh, R., Surachawala, T. & Lerman, K. Entropy-based Classification of Retweeting Activity on Twitter. In *Proc. KDD workshop on Social Network Analysis*, SNAKDD'11, (ACM, New York, NY, USA, 2011).
40. Armstrong, J. S. & Collopy, F. Error measures for generalizing about forecasting methods: Empirical comparisons. *Int. J. Forecasting* **8**, 69–80 (1992).

Acknowledgments

We thank Tawan Surachawala and Suradej Intagorn for their help with data collection. This work was supported in part by AFOSR (contract FA9550-10-1-0569), by NSF (grant CIF-1217605) and by DARPA (contract W911NF-12-1-0034).

Author contributions

N.H. and K.L. developed the model. N.H. performed empirical analysis and evaluation. Both authors contributed to the manuscript.

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/scientificreports>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Hodas, N.O. & Lerman, K. The Simple Rules of Social Contagion. *Sci. Rep.* **4**, 4343; DOI:10.1038/srep04343 (2014).

This work is licensed under a Creative Commons Attribution 3.0 Unported license. To view a copy of this license, visit <http://creativecommons.org/licenses/by/3.0>