

Social dynamics of Digg

Tad Hogg¹ and Kristina Lerman^{2*}

*Correspondence: lerman@isi.edu

²USC Information Sciences Institute,
Marina del Rey, CA, USA
Full list of author information is
available at the end of the article

Abstract

Online social media provide multiple ways to find interesting content. One important method is highlighting content recommended by user's friends. We examine this process on one such site, the news aggregator Digg. With a stochastic model of user behavior, we distinguish the effects of the content visibility and interestingness to users. We find a wide range of interest and distinguish stories primarily of interest to a users' friends from those of interest to the entire user community. We show how this model predicts a story's eventual popularity from users' early reactions to it, and estimate the prediction reliability. This modeling framework can help evaluate alternative design choices for displaying content on the site.

1 Introduction

The explosive growth of the Social Web hints at collective problem-solving made possible when people have tools to connect, create and organize information on a massive scale. The social news aggregator Digg, for example, allows people to collectively identify interesting news stories. The microblogging service Twitter has created a cottage industry of third-party applications, such as identifying trends from the millions of conversations taking place on the site and notifying you when your friends are nearby. Other sites enable people to collectively create encyclopedias, develop software, and invest in social causes. Analyzing records of complex social activity can help identify communities and important individuals within them, suggest relevant readings, and identify events and trends.

Effective use of this technology requires understanding how the social dynamics emerges from the decisions made by interconnected individuals. One approach is a stochastic modeling framework, which represents each user as a stochastic process with a few states. As an example, applying this approach to Digg successfully described observed voting patterns of Digg's users [1–3]. However, quantitative evaluation of the model was limited by the poor quality of data, which was extracted by scraping Digg's web pages.

In this paper we present two refinements to this modeling approach for Digg. First, we explicitly allow for systematic differences in *interest* in news stories for linked and unlinked users. This distinction is a key aspect of social media where links indicate commonality of user interests. We also include additional aspects of the Digg user interface in the model, thereby accounting for cases where the existing model identified anomalous behaviors. As the second major contribution, we describe how to measure confidence intervals of model predictions. We show that confidence intervals are highly correlated with the error between the predicted and actual votes stories accrue. Thus the confidence intervals indicate the quality of the model's predictions on a per-user or per-story basis.

This paper is organized as follows. The next section describes Digg and our data set. We then present a stochastic model of user behavior on Digg that explicitly includes dependencies on social network links. Using this model, we quantify these dependencies and discuss how the model predicts eventual popularity of newly submitted content. Finally, we compare our approach with other studies and discuss possible applications of stochastic models incorporating social network structure.

2 Digg: a social news portal

At the time data was collected, Digg was a popular news portal with over 3 million registered users. Digg allowed users to submit and rate news stories by voting on, or ‘digging,’ them. Every day Digg promoted a small fraction of submitted stories to the highly visible *front page*. Although the exact promotion mechanism was kept secret and changes occasionally, it appears to use the number of votes the story receives. Digg’s popularity was largely due to the emergent front page created by the collective decisions of its many users. Below we describe the user interface that existed at the time of data collection.

2.1 User interface

Submitted stories appear in the *upcoming* stories list, where they remain for about 24 hours or until promoted to the front page. By default, Digg shows upcoming and front page stories in recency lists i.e., in reverse chronological order with the most recently submitted (promoted) story at the top of the list. A user may choose to display stories by popularity or by some broad topic. Popularity lists show stories with the most votes up to that time, e.g., the most popular stories submitted (promoted) in the past day or week. Each list is divided into pages, with 15 stories on each page, and the user has to click to see subsequent pages.

Digg allows users to designate friends and track their activities. The friend relationship is asymmetric. When user *A* lists user *B* as a *friend*, *A* can follow the activities of *B* but not vice versa. We call *A* the *fan*, or follower, of *B*. The *friends interface* shows users the stories their friends recently submitted or voted for.^a

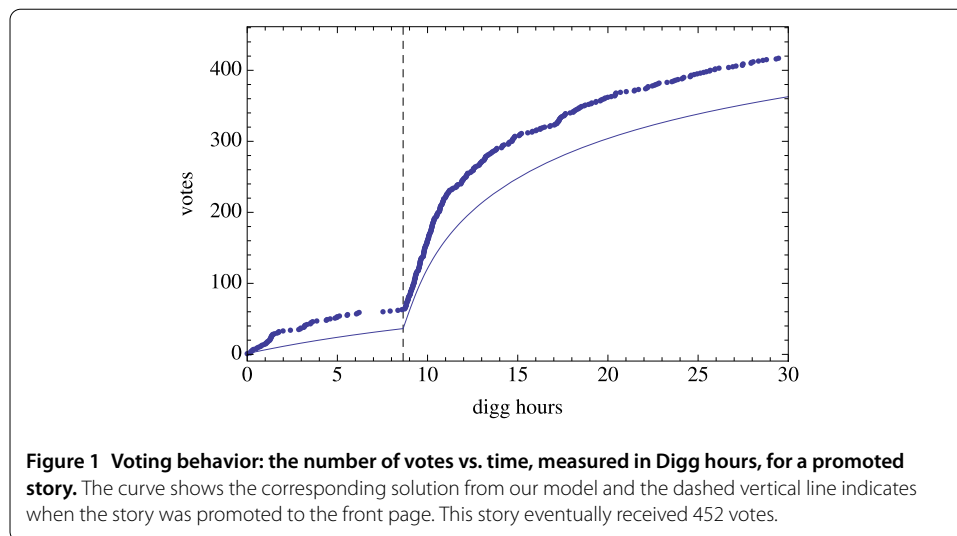
In this paper, we focus on the recency and ‘popular in the last 24 hours’ lists for all stories and the friends interface list for each user. These lists appear to account for most of the votes a story receives.

2.2 Evolution of story popularity

Most Digg users focus on front page stories, so upcoming stories accrue votes slowly. When a story is promoted to the front page, it becomes visible to many more users and accrues votes rapidly. Figure 1 shows the evolution of the number of votes for a story submitted in June 2009. The slope abruptly increases at promotion time (dashed line). As the story ages, accumulation of new votes slows down, and after a few days stories typically no longer receive new votes.

The final number of votes varies widely among the stories. Some promoted stories accumulate thousands of votes, while others muster only a few hundred. Stories that are never promoted receive few votes, in many cases just a single vote from the submitter, and are removed after about 24 hours.

A challenge for understanding this variation in popularity is the interaction between the stories’ *visibility* (how Digg displays them) and their *interestingness* to users. Models



accounting for the structure of the Digg interface can help distinguish these contributions to story popularity.

2.3 Data set

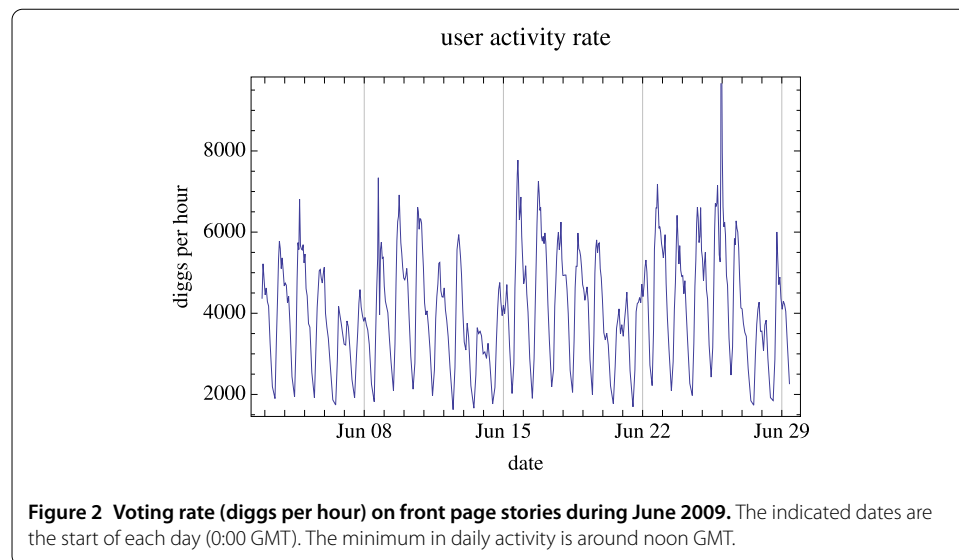
We used Digg API to collect complete (as of 2 July 2009) voting histories of all stories promoted to the front page of Digg in June 2009.^b For each story, we collected story id, submitter's id, and the list of voters with the time of each vote. We also collected the time each story was promoted to the front page. The data set contains over 3 million votes on 3,553 promoted stories. We did not retrieve data about stories that were submitted to Digg during that time period but were never promoted. Thus, in contrast to prior models that included promotion behavior [2], our focus is on the behavior of promoted stories, which receive most of the attention from Digg users.

We define an *active user* as any user who voted for at least one story on Digg during the data collection period. Of the 139,409 active users, 71,367 designated at least one other user as a friend. We extracted the friends of these users and reconstructed the fan network of active users, i.e., a directed graph of active users who are following activities of other users.

Over the period of a month, some of the voters in our sample deleted their accounts and were marked 'inactive' by Digg. Such cases represent a tiny fraction of all users in the data set; therefore, we take the number of users to be constant.

2.4 Daily activity variation

Activity on Digg varies considerably over the course of a day, as seen in Figure 2. Adjusting times by the cumulative activity on the site accounts for this variation and improves predictions. Following [4, 5] we define the 'Digg time' between two events (e.g., votes on a story) as the total number of votes made during the time between those events. With our data, this only counts votes on stories that were eventually promoted to the front page. In our data set, there are on average about 4,000 votes on front page stories per hour, with a range of about a factor of 3 in this rate during the course of a day. This behavior is similar to that seen in an extensive study of front page activity in 2007 [4], and as in that study we



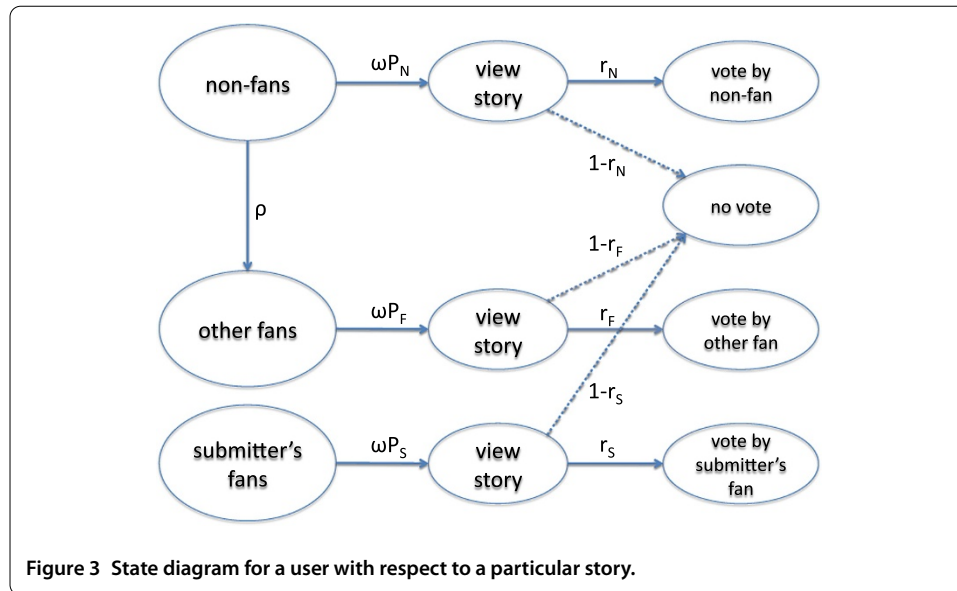
scale the measure by defining a ‘Digg hour’ to be the average number of front page votes in an hour.

3 Social dynamics of Digg

A key challenge in stochastic modeling is finding a useful combination of simplicity, accuracy and available data to calibrate the model. Stochastic models of online social media describe the joint behavior of their users and content. Since these web sites receive much more content than users have time or interest to examine, one important property to model is how readily users can find content. A second key property is how users react to content once they find it. Thus an important modeling choice for social media is the level of detail sufficient to distinguish user behavior and content visibility. Following the practice of population dynamics [6] and epidemic modeling [7] we consider groups of users and content. We assume that individuals within each group have sufficiently similar behavior that their differences do not affect the main questions of interest. In the case of Digg, one such question is the number of votes a story receives over time. In our approach, we focus on how a single story accumulates votes, based on the combination of how easily users can find the story and how interesting it is to different groups of users.

Following Refs. [1, 2], we start with a simple model in which story visibility is determined primarily by its location on the recency and friends lists, and use a single value to describe the story’s interestingness to the user community. We use the ‘law of surfing’ [8] to relate location of the story to how readily users find it. This model successfully captured the qualitative behavior of typical stories on Digg and how that behavior depended on the number of fans of the story’s submitter [2, 9].

However, the simple model did not quantitatively account for several behaviors in the new data set. These included the significant daily variation in activity rates seen in Figure 2 and systematic differences in behavior between fans of a story’s submitter and other users. In particular, the new data was sufficiently detailed to show users tend to find stories their friends submit as more interesting than stories friends vote on but did not submit. Another issue with the earlier model is a fairly large number of votes on stories far down the recency list. This is especially relevant for upcoming stories where the large rate of



new submissions means a given story remains near the top of the recency list for only a few minutes. To account for such votes, the model's estimated 'law of surfing' parameters indicated users browse an implausibly large number of pages while visiting Digg.


These observations motivate the more elaborate model described in this paper. This model includes systematic differences in interestingness between fans and other users and additional ways Digg makes stories visible to users.

3.1 User model

We allow for differences between users by separating them into groups, and assume that visibility of stories and voting behavior of users within each group is statistically the same. We refine the previous stochastic model of Digg [2] by not only distinguishing votes from fans and non-fans [5], but also allowing for differences between fans of the *submitter* and fans of *other* voters who are not also fans of the submitter. The state diagram Figure 3 shows the resulting user model. The state *submitter's fans* includes all users who are fans of the submitter and have not yet seen the story; the *other fans* state includes all users who are fans of other voters but not the submitter, who have not yet seen the story; and the *non-fans* state includes all users who are neither fans of the submitter nor other voters, and have not yet seen the story. The state *no vote* includes all users who have seen the story and decided not to vote for it. With respect to votes on a given story treated in this model, users visit Digg according to a Poisson process with average rate ω in terms of Digg time.

Users transition between states stochastically by browsing Digg's web pages and voting for stories. The submitter provides a story's first vote. All of her fans start in the *submitter's fans* state, and all other users start in the *non-fans* state. Each vote causes non-fan users who are that voter's fans and who have not yet seen the story to transition from the *non-fans* state to the *other fans* state. A user making this transition is not aware of that change until later visiting Digg and seeing the story on her friends list.

Once a user sees a story, she will vote for it with probability given by how interesting she finds the story. Nominally people become fans of those whose contributions they consider



interesting, suggesting fans have a systematically higher interest in stories than non-fans. Our model accounts for this by having the probability a user votes on a story depend on the user's state. Users in each state also have a different probability to see stories, which is determined by the story's *visibility* to that category of users. Users vote at most once on a story, and our focus is on the final decision to vote or not after the user sees the story.

The visibility of stories changes as stories age and accrue votes. Figure 4 shows the state diagram of stories. A story starts at the top of the upcoming stories recency list. The location increases with each new submission. A promoted story starts at the top of the front pages. The location increases as additional stories are promoted.

These state diagrams lead to a description of the average rates of growth [3] for votes from submitter fans, other fans and non-fans of prior voters, v_S , v_F and v_N , respectively:

$$\frac{dv_S}{dt} = \omega r_S P_S S, \quad (1)$$

$$\frac{dv_F}{dt} = \omega r_F P_F F, \quad (2)$$

$$\frac{dv_N}{dt} = \omega r_N P_N N, \quad (3)$$

where t is the Digg time since the story's submission and ω is the average rate a user visits Digg (measured as a rate per unit Digg time). We find only a small correlation between voting activity and the number of fans. Thus we use the average rate users visit Digg, rather than having the rate depend on the number of fans a user has. v_N includes the vote by the story's submitter. P_S , P_F and P_N denote the story's *visibility* and r_S , r_F and r_N denote the story's *interestingness* to users who are submitter fans, other fans or non-fans of prior voters, respectively. Visibility depends on the story's state (e.g., whether it has been promoted), as discussed below. Interestingness is the probability a user who sees the story will vote on it.

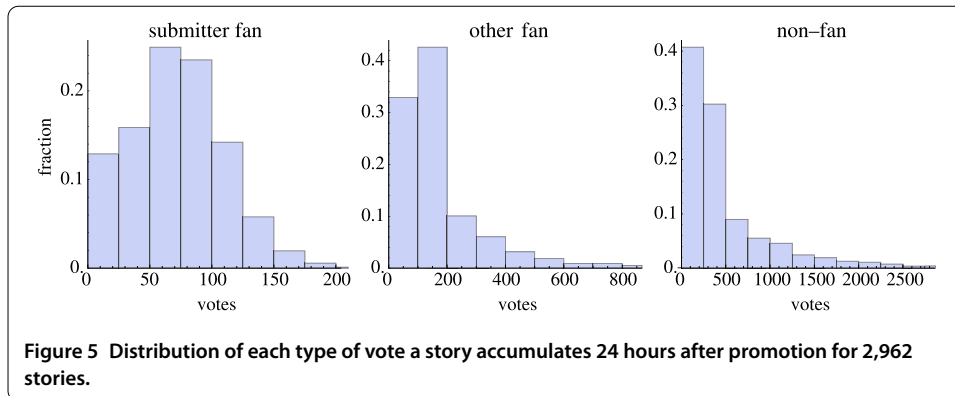
These voting rates depend on the number of users in each category who have not yet seen the story: S , F and N . The quantities change as users see and vote on the story, with average rate of change given by

$$\frac{dS}{dt} = -\omega P_S S, \quad (4)$$

$$\frac{dF}{dt} = -\omega P_F F + \rho N \frac{dv}{dt}, \quad (5)$$

$$\frac{dN}{dt} = -\omega P_N N - \rho N \frac{dv}{dt}, \quad (6)$$

with $v = v_S + v_F + v_N$ the total number of votes the story has received. The quantity ρ is the probability a user is a fan of the most recent voter, conditioned on that user not having



seen the story nor being a fan of any voter prior to the most recent voter. For simplicity, we treat this probability as a constant over the voters, thus averaging over the variation due to clustering in the social network and the number of fans a user has. The first term in each of these equations is the rate the users see the story. The second terms arise from the rate the story becomes visible in the friends interface of users who are not fans of previous voters but are fans of the most recent voter.

Initially, the story has one vote (from the submitter) and the submitter has S_0 fans, so $v_S(0) = v_F(0) = 0$, $v_N(0) = 1$, $S(0) = S_0$, $F(0) = 0$ and $N(0) = U - S_0 - 1$ where U is the total number of active users at the time the story is submitted. Over time, a story becomes less visible to users as it moves down the upcoming or (if promoted) front page recency lists, thereby attracting fewer votes and hence fewer new fans of prior voters. If the story gathers more votes than other stories, it is moved to the top of the popularity list, so becomes more visible.

Figure 5 shows the range of votes the stories receive by 24 hours after promotion. Generally, stories have most votes from non-fans, somewhat fewer from other fans and a relatively small number from submitter's fans. The number of votes from submitter's fans is weakly correlated with the numbers from other fans (correlation coefficient 0.11) and non-fans (0.12). The numbers from other fans and non-fans are highly correlated (0.95).

3.2 Story visibility

A fan easily sees the story via the friends interface, so we take $P_S = P_F = 1$ for front page stories. While the story is upcoming, it appears in the friends interface but users do not necessarily choose to view upcoming stories friends liked. Users can readily make this choice because the friends interface distinguishes upcoming from front page stories. We characterize the lower visibility of upcoming stories with constants c_S and c_F which are less than 1. The corresponding visibility is then $P_S = c_S$ and $P_F = c_F$.

Users who are not fans of prior voters must find the story on the front or upcoming pages. Thus P_N depends on how users navigate through these pages and the story's location at the time the user visits Digg. This navigation is not given by our data. Instead, we use a model of user navigation through a series of web pages in which users estimate the value of continuing at the site, and leave when that value becomes negative [8]. This 'law of surfing' leads to an inverse Gaussian distribution of the number of pages m a user visits

before leaving the web site,

$$e^{-\lambda(m-\mu)^2/(2m\mu^2)} \sqrt{\frac{\lambda}{2\pi m^3}} \quad (7)$$

with mean μ and variance μ^3/λ [8]. We use this distribution for user navigation on Digg [2].

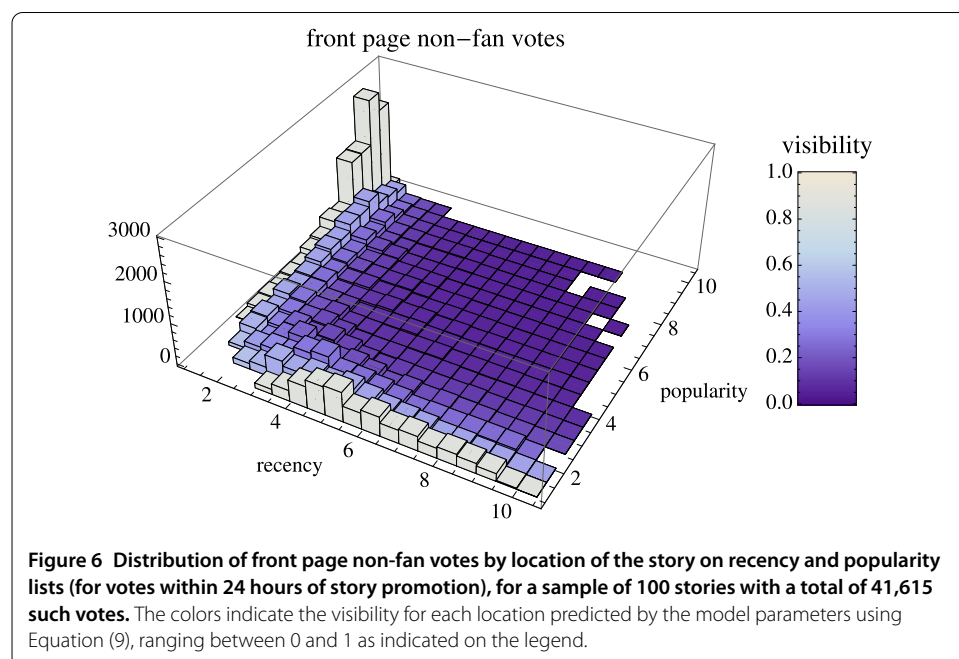
The visibility of a story on the m th front or upcoming page is the fraction of users who visit *at least* m pages, i.e., the upper cumulative distribution of Equation (7). For $m > 1$, this fraction is

$$f_{\text{page}}(m) = \frac{1}{2} (F_m(-\mu) - e^{2\lambda/\mu} F_m(\mu)), \quad (8)$$

where $F_m(x) = \text{erfc}(\alpha_m(m-1+x)/\mu)$, erfc is the complementary error function, and $\alpha_m = \sqrt{\lambda/(2(m-1))}$. For $m = 1, f_{\text{page}}(1) = 1$. The visibility of stories decreases in two distinct ways when a new story arrives. First, a story moves down the list on its current page. Second, a story at the 15th position moves to the top of the next page. For simplicity, we model these processes as decreasing visibility in the same way through m taking on fractional values within a page, e.g., $m = 1.5$ denotes the position of a story half way down the list on the first page.

Digg presents several lists of stories. We focus on two lists as the major determinants of visibility for front page stories: reverse chronological order ('recency') and most popular in the past 24 hours ('popularity'). Users can also find stories via other means. For instance, Digg includes other lists showing recent and popular stories in specific topics (e.g., sports or business) and popularity over longer time periods, e.g., the previous week. Stories on Digg may also be linked to from external web sites (e.g., the submitter's blog).

For front page votes, the recency and popularity lists provide the bulk of non-fan votes while the stories are close to the top of at least one of these lists, as illustrated in Figure 6.



The rank on the recency list is the number of stories promoted more recently than that story and the rank on the popularity list is the number of stories, promoted within the 24 hours prior to the vote, with more votes. Since a page shows 15 stories, the location in terms of number of pages, as shown in the figure, is 1/15th the rank, starting from page 1. Some votes occur while stories are far down both the recency and popularity lists, so the user likely finds the story by another method.

From these observations, we account for three ways users who are not connected to prior voters find a story: via the recency list, via the popularity list or via one of the other methods described above. We combine visibility from these three methods assuming independent choices by users, giving the probability to see the story as

$$P_{\text{visibility}}(t, v) = 1 - (1 - f_{\text{page}}(p(t)))(1 - f_{\text{page}}(p_{\text{popularity}}(v)))(1 - \beta), \quad (9)$$

where $p(t)$ and $p_{\text{popularity}}(v)$ are the locations of the story on the recency and popularity lists, respectively, and β is the probability to find the story by another method. Although the positions of the stories on these lists depend on the specific stories submitted or promoted shortly after the story, these locations are approximately determined by the time t and number of votes v the story has, as described below. For visibility by other methods, we take β to be a constant, independent of story properties such as time since submission or number of votes. That is, we do not explicitly model factors affecting the visibility of stories by other methods. Moreover, the data does not provide information on how users find a story nor the number who find the story but choose not to vote on it. These data limitations preclude a more detailed model of these other methods by which users find stories. For our focus on promoted stories, this is a minor limitation because the recency and popularity lists account for the bulk of the non-fan votes determined by our parameter estimates discussed below.

The location of a story on the recency and popularity lists could be additional state variables, which change as new stories are added and gain votes. Instead of modeling this in detail, we approximate these locations using the close relation between location and time (for recency) or votes (for popularity).

Position of a story in the recency list

Using Digg time to account for the daily variation in activity on the site, the rate of story submission and promotion is close to linear. Thus the page number of a story on either the upcoming or front page recency list is [2]

$$p(t) = \begin{cases} k_u t + 1 & \text{if } t < T_{\text{promotion}}, \\ k_f(t - T_{\text{promotion}}) + 1 & \text{otherwise,} \end{cases} \quad (10)$$

where $T_{\text{promotion}}$ is the time the story is promoted to the front page and the slopes are given in Table 1. Since each page holds 15 stories, these rates are 1/15th the story submission and promotion rates, respectively.

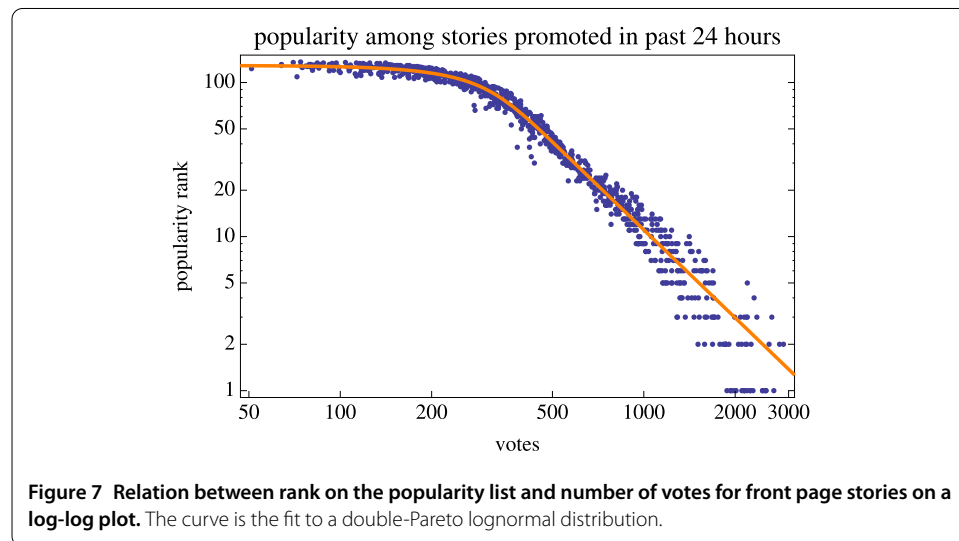
Position of a story in the popularity list

The position of a story on the popularity list is the number of stories submitted or promoted in the previous 24 hours with more votes, for stories on upcoming or front page

Table 1 Model parameters, with times in Digg hours

Parameter	Value
average rate each user visits Digg	$\omega = 0.16 \pm 0.01/\text{hr}$
number of active users	$U = 248,000 \pm 3,000$
page view distribution	$\mu = 0.92 \pm 0.04$ $\lambda = 0.9 \pm 0.1$
visibility by other methods	$\beta = 0.05 \pm 0.01$
probability a user is a voter's fan	$\rho = 1.7 \times 10^{-5}$
upcoming stories location	$k_u = 59.8 \text{ pages/hr}$
front page location	$k_f = 0.31 \text{ pages/hr}$
fraction viewing upcoming pages	
submitter fans	$c_S = 0.57 \pm 0.03$
other fans	$c_F = 0.10 \pm 0.01$
non-fans	$c_N = 0.11 \pm 0.01$
story specific parameters	
interestingness to submitter fans	r_S
interestingness to other fans	r_F
interestingness to non-fans	r_N
number of submitter's fans	S_0
promotion time	$T_{\text{promotion}}$

Error ranges are 95% confidence intervals for the parameter estimates.



lists, respectively. The distribution of votes among stories in a 24 hour period is similar from day to day. Thus we expect that a story's position on the popularity list, determined by the location of its number of votes in this distribution, is approximately a function of its number of votes alone, with only minor variation depending on the time (i.e., the set of other stories from the past 24 hours). To evaluate this expectation, we determined the popularity rank of a sample of front page votes within 24-hours of story promotion and compared the rank to the number of votes, as shown in Figure 7. As we expected, the figure shows that the rank is well-determined by the number of votes. Thus we model the position of a story on the popularity list as depending only on the number of votes the story has at the time, i.e., consider $p_{\text{popularity}}(v)$ as a function of the number of votes the story has. This gives a simple, approximate relation between the actual location and number

of votes - ignoring the minor variations due to the specific stories promoted at different times.

To identify a suitable functional approximation for $p_{\text{popularity}}(v)$, we note that a story typically accumulates votes at a rate proportional to how interesting it is to the user population. From a prior analysis of votes in 2006 on Digg [2] we expect the interestingness to be approximately lognormally distributed. Thus if we observe a sample of votes on stories over the same time interval for each story, the distribution of votes, and hence location on the popularity list, would follow a lognormal distribution. However, the popularity list includes stories of various times up to 24 hours since submission or promotion. Thus some stories of high interest will have few votes because they were just recently submitted or promoted, and conversely some stories with only moderate interestingness will have relatively many votes because they have been available for votes for nearly 24 hours. The combination of lognormal distribution of rates for accumulating votes and this variation in the observation times modifies the tails of the lognormal to be power-law, i.e., a double-Pareto lognormal distribution [10].

Such a distribution fits the observed positions on the front page popularity list, as indicated in Figure 7. The fit for the rank (i.e., number of stories above the given one in the popularity list, so the story promoted in the past 24 hours with the most votes has rank 0) is

$$\text{rank} = S(1 - \Lambda(a, b, v, \sigma; v)), \quad (11)$$

where $S = 129.0 \pm 0.1$ is the average number of stories promoted in 24 hours and $\Lambda(\dots; v)$ is the cumulative distribution of a double-Pareto lognormal distribution, i.e., fraction of cases with fewer than v votes. The parameters $a = 1.90 \pm 0.005$ and $b = 2.50 \pm 0.03$ are the power-law exponents for the upper and lower tails of the distribution, respectively, and the parameters $\nu = 5.88 \pm 0.002$ and $\sigma = 0.16 \pm 0.004$ characterize the location and spread of the lognormal behavior in the center of the distribution. This fit captures the power-law tail relating stories near the top of the popularity list with the number of votes the story has. These are the cases for which the popularity list contributes significantly to the overall visibility of a story. More precisely, the Kolmogorov-Smirnov (KS) statistic shows the vote counts are consistent with this distribution (p -value 0.92). We use this distribution, combined with the rate stories are promoted, to relate the number of votes a story has to its position on the popularity list, providing a functional form for $p_{\text{popularity}}(v)$.

The popularity rank for upcoming stories submitted in the past 24 hours is more complicated than for the front pages due to the promotion. Nevertheless, the main factor determining a story's location on the popularity list is its number of votes, particularly after adjusting for most of the daily variation in activity by using Digg time described in Section 2.4. Stories with many votes are more likely to be promoted, and hence removed from the popularity list for upcoming stories. This removal alters the upper tail of the distribution and hence the numbers of votes for stories appearing near the top of the popularity list. Moreover, out of the $\approx 20,000$ stories submitted each day, our data includes only the ≈ 100 stories per day eventually promoted. However, popularity significantly contributes to visibility only for stories near the top of the popularity list. Thus for our model, it is sufficient to determine the relation between votes and rank for upcoming stories with relatively many votes. Such stories are likely to be promoted eventually and hence included

in our sample. Instead of a power-law tail, our data on the eventually promoted stories is better fit by an exponential for the upcoming stories with relatively many votes, and hence near the top of the popularity list:

$$\text{rank} = e^{c-dv} \quad (12)$$

with $c = 5.3 \pm 0.01$ and $d = 0.029 \pm 0.0002$. This fits well for upcoming stories submitted within the past 24 hours with more than 100 votes, corresponding to rank of about 20 or less on the upcoming popularity list. For stories with few votes, e.g., fewer than 10 or 20, this fit based on the stories eventually promoted substantially underestimates the rank. Nevertheless, the estimated rank for such stories is still sufficiently large that the law of surfing parameters we estimate indicate users do not find such stories via the popularity list. Thus this underestimate does not significantly affect our model's behavior for upcoming stories.

Friends interface

The fans of the story's submitter can find the story via the friends interface. As additional people vote on the story, their fans can also see the story. We model this with $F(t)$, the number of fans of voters on the story by time t who are not also fans of the submitter and have not yet seen the story. Although the number of fans is highly variable, we use the average number of additional fans from an extra vote, ρN , in Equation (4).

4 Parameter estimation

We estimate model parameters using 100 stories from the middle of our sample.

4.1 Estimating parameters from observed votes

In our model, story location affects visibility only for non-fan voters since fans of prior voters see the story via the friends interface. Thus we use just the non-fan votes to estimate visibility parameters, via maximum likelihood. Specifically, we use the non-fan votes to estimate the 'law of surfing' parameters μ and λ , as well as the probability for finding the story some other way, β .

This estimation involves comparing the observed votes to the voting rate from the model. As described above, the model uses rate equations to determine the average behavior of the number of votes. We relate this average to the observed number of votes by assuming the votes from non-fan users form a Poisson process whose expected value is $dv_N(t)/dt$, given by Equation (3).

For a Poisson process with a constant rate ν , the probability to observe n events in time T is the Poisson distribution $e^{-\nu T}(\nu T)^n/n!$. This probability depends only on the *number* of events, not the specific times at which they occur. Estimating ν involves maximizing this expression, giving $\nu = n/T$. Thus the maximum-likelihood estimate of the rate for a constant Poisson process is the average rate of the observed events.

In our case, the voting rate changes with time, requiring a generalization of this estimation. Specifically consider a Poisson process with nonnegative rate $\nu(t)$ which depends on one or more parameters to be estimated. Thus in a small time interval $(t, t + \Delta t)$, the probability for a vote is $\nu(t)\Delta t$, and this is independent of votes in other time intervals, by the definition of a Poisson process. Suppose we observe n votes at times $0 < t_1 < t_2 < \dots < t_n < T$

during an observation time interval $(0, T)$. Considering small time intervals Δt around each observation, the probability of this observation is

$$\begin{aligned} & P(\text{no vote in } (0, t_1))v(t_1)\Delta t \\ & \times P(\text{no vote in } (t_1, t_2))v(t_2)\Delta t \\ & \dots \\ & \times P(\text{no vote in } (t_{n-1}, t_n))v(t_n)\Delta t \\ & \times P(\text{no vote in } (t_n, T)). \end{aligned}$$

The probability for no vote in the interval (a, b) is

$$\exp\left(-\int_a^b v(t) dt\right).$$

Thus the log-likelihood for the observed sequence of votes is

$$-\int_0^T v(t) dt + \sum_i \log v(t_i). \quad (13)$$

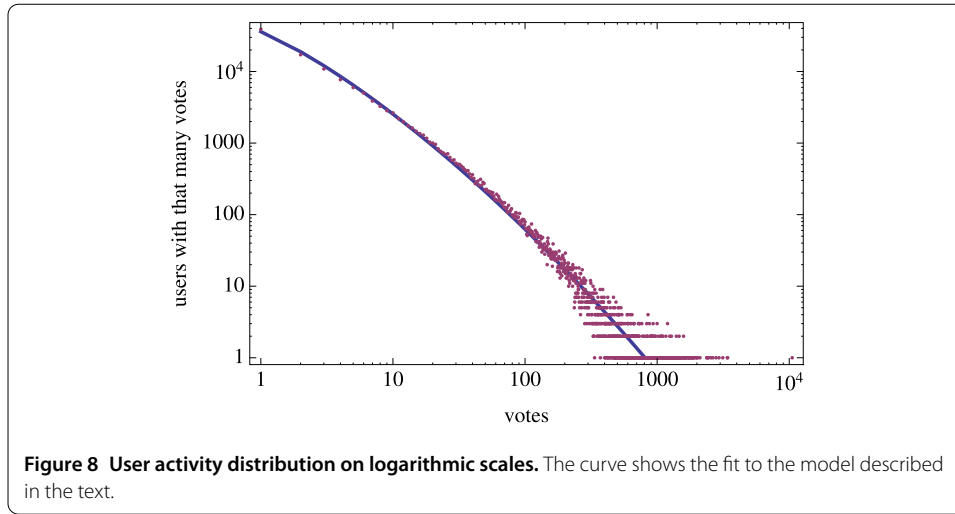
The maximum-likelihood estimation for parameters determining the rate $v(t)$ is a trade-off between these two terms: minimizing $v(t)$ over the range $(0, T)$ to increase the first term while maximizing the values $v(t_i)$ at the times of the observed votes. If $v(t)$ is constant, this likelihood expression simplifies to $-vT + n \log v$ with maximum at $v = n/T$ as discussed above for the constant Poisson process. When $v(t)$ varies with time, the maximization selects parameters giving relatively larger $v(t)$ values where the observed votes are clustered in time.

We combine this log-likelihood expression from the votes on several stories, and maximize the combined expression with respect to the story-independent parameters of the model, while determining the interestingness parameters separately for each story.

4.2 User activity

Our model involves a population of active users who visit Digg during our sample period and vote on stories. Specifically, the model uses the rate users visit Digg, ωU . We do not observe visits in our data, but can infer the relevant number of active users, U , from the heterogeneity in the number of votes by users. The data set consists of 139,409 users who voted at least once during the sample period, giving a total of 3,018,197 votes. Figure 8 shows the distribution of this activity. Most users have little activity during the sample period, suggesting a large fraction of users vote infrequently enough to never have voted during the time of our data sample. This behavior can be characterized by an activity rate for each user. A user with activity rate v will, on average, vote on vT stories during a sample time T . We model the observed votes as arising from a Poisson process whose expected value is vT and the heterogeneity arising from a lognormal distribution of user activity rates [11]:

$$P_{\text{lognormal}}(\mu, \sigma; r) = \frac{1}{\sqrt{2\pi}r\sigma} \exp\left(-\frac{(\mu - \log(r))^2}{2\sigma^2}\right), \quad (14)$$



where parameters μ and σ are the mean and standard deviation of $\log(r)$.

This model gives rise to the extended activity distribution while accounting for the discrete nature of the observations. The latter is important for the majority of users, who have low activity rates and vote only a few times, or not at all, during our sample period.

Specifically, for n_k users with k votes during the sample period, this mixture of lognormal and Poisson distributions [12, 13] gives the log-likelihood of the observations as

$$\sum_k n_k \log P(\mu, \sigma; k),$$

where $P(\mu, \sigma; k)$ is the probability of a Poisson distribution to give k votes when its mean is chosen from a lognormal distribution $P_{\text{lognormal}}$ with parameters μ and σ . From Equation (14),

$$P(\mu, \sigma; k) = \frac{1}{\sqrt{2\pi} \sigma k!} \int_0^\infty \rho^{k-1} e^{-(\log(\rho) - \mu)^2 / (2\sigma^2) - \rho} d\rho$$

for integer $k \geq 0$. We evaluate this integral numerically. In terms of our model parameters, the value of μ in this distribution equals νT .

Since we do not observe the number of users who did not vote during our sample period, i.e., the value of n_0 , we cannot maximize this log-likelihood expression directly. Instead, we use a zero-truncated maximum likelihood estimate [14] to determine the parameters μ and σ for the vote distribution of Figure 8. Specifically, the fit is to the probability of observing k votes conditioned on observing at least one vote. This conditional distribution is $P(\mu, \sigma; k) / (1 - P(\mu, \sigma; 0))$ for $k > 0$, and the corresponding log-likelihood is

$$\sum_{k>0} n_k \log P(\mu, \sigma; k) - U_+ \log(1 - P(\mu, \sigma; 0)),$$

where U_+ is the number of users with at least one vote in our sample. Maximizing this expression with respect to the distribution's parameters μ and σ gives νT lognormally distributed with the mean and standard deviation of $\log(\nu T)$ equal to -0.10 ± 0.04 and

2.43 ± 0.02 , respectively. Based on this fit, the curve in Figure 8 shows the expected number of users with each number of votes. This is a discrete distribution: the lines in the figure between the expected values serve only to distinguish the model fit from the points showing the observed values.

With these estimated parameters, $P(\mu, \sigma; 0) = 0.43$, indicating 43% of the users had sufficiently low, but nonzero, activity rate that they did not vote during the sample period. We use this value to estimate U , the number of active users during our sample period: $U = U_+ / (1 - P(\mu, \sigma; 0))$.

4.3 Links among users

We observe $u = 258,218$ users with fans, and these users have a total of $c = 1,731,658$ connections. Our data has 139,409 distinct voters, of which 78,007 have no fans. There is little correlation between links and voting activity, so we estimate the fraction of users with zero fans from the ratio of these values, i.e., about 56%. Thus the average number of fans per user, including users without fans, is $c/(1.56u) \approx 4.3$.

We estimate the model parameter ρ of Equations (4) and (6) as the probability a fan link connects the first to the second user of a randomly selected pair of users, corresponding to the average number of fans per user divided by the number of active users U .

4.4 Visibility to submitter's fans

Because stories are always visible to fans and we know the number of fans of the story's submitter, the model behavior (Equations (1) and (4)) can be solved without reference to the rest of the model. We have $P_S = 1$ when the story is on the front page and $P_S = c_S < 1$, reflecting users' preference for front page stories. Thus, these equations have two story-independent parameters, i.e., the rate users visit Digg (ω) and the probability users view upcoming stories submitted by their friends (c_S), and two story-dependent parameters, i.e., the interestingness (r_S) and number of fans of the submitter (S_0). S_0 is given in our data, while we estimate the other parameters from the data, i.e., votes by fans of the stories' submitters.

4.5 Visibility to non-fans

In our model, story location affects visibility only for non-fan voters since fans of prior voters see the story via the friends interface. Thus we use just the non-fan votes to estimate visibility parameters, via maximum likelihood. A story typically receives only a few dozen votes before promotion, mostly from fans. With the value of ρ , estimated as described above, Equation (6) gives $N(t) \approx U$ up to a few hours after promotion. Over this time period, Equation (3) simplifies to $dN/dt \approx \omega U r_N P_N$ with P_N depending on story location on the recency and popularity lists. r_N is constant for a given story, so P_N determines the time variation in the voting rate by non-fans.

For front page stories, in our model $P_N = P_{\text{visibility}}(t, \nu)$ from Equation (9), which has three parameters: μ and λ characterizing the browsing behavior for the recency and popularity lists, and the probability to find the story by other methods, β . We estimate these parameters by maximizing the likelihood of observing the non-fan front page votes according to the model, as described above for estimating a Poisson process with a time-dependent rate in Equation (13). This estimation also determines r_N for each story.

For upcoming stories, we take $P_N = c_N P_{\text{visibility}}(t, \nu)$, giving a single additional parameter, c_N , to estimate, since we assume browsing behavior on the upcoming pages is the same as

for front pages. This assumption has little effect on the model behavior because of the large number of submissions and relatively few non-fan votes for upcoming stories. A submitted story remains near the front of the recency list for only about a minute after submission and stories reaching the front of the popularity list (due to having many votes) are soon promoted to the front page. Thus moderate variations in how deeply users browse the upcoming recency or popularity lists (i.e., the values for μ and λ) have little effect on the non-fan votes. Instead, the relatively few non-fan upcoming votes arise mainly through users finding the story by other means. That is, in most cases $P_{\text{visibility}} \approx \beta$ for upcoming stories. Thus $P_N = c_N P_{\text{visibility}}(t, v) \approx c_N \beta$ and any difference between β for upcoming and front page stories would merely rescale the value of c_N . This parameter is readily estimated using Equation (13) with the upcoming non-fan votes.

4.6 Visibility to other fans

From Equation (2), dv_F/dt changes abruptly when the story is promoted since P_F changes from c_F to 1 upon promotion. Thus we estimate c_F by the change in voting rate by fans other than those of the submitter by comparing the votes a story receives one hour before promotion and the votes received during the hour after promotion.

With all story-independent parameters estimated, we can then solve the full model for a story to determine dv_F/dt as a function of time. This gives the expected rate of other fan votes as a function of time. We determine r_F for the story as the value maximizing the log-likelihood (Equation (13)) for the other fan votes the story receives.

4.7 Summary

Table 1 lists the estimated parameters. All of these parameters, except the three story interestingness parameters r_S , r_F and r_N , are either known (e.g., the number of submitter's fans) or estimated from data from a sample of stories and then used for all stories. The interestingness parameters are estimated individually for each story from its votes.

5 Results

Figure 1 compares the solution of the rate equations with the actual votes for one story. This illustrates that the model captures the main qualitative features of the vote dynamics: an abrupt jump in votes after promotion followed by a slowing of the voting rate.

Figure 6 shows how visibility estimated by our model (indicated by color) compares with the distribution of front page votes. Many votes occur when the story is recently promoted (so near the top of the recency list) or has received many votes within 24 hours after promotion (so near the top of the popularity list). This is consistent with our model, which predicts higher visibility for stories in these positions on the lists.

5.1 Interestingness for fans and non-fans

We use the model to evaluate systematic differences in story interestingness between fans and non-fans. The estimated r values indicate the stories have a wide range of interestingness to users, as shown in Figure 9, along with fits to lognormal distributions. The figure shows r_N values tend to be much smaller than the interestingness for fans, as also seen in an earlier study with a smaller data set from 2006 [5]. The r -values are weakly correlated, with Spearman rank correlation between r_S and r_F of 0.17, between r_S and r_N of 0.03, and between r_F and r_N of 0.39. Moreover, there is a large range in the ratio of interestingness

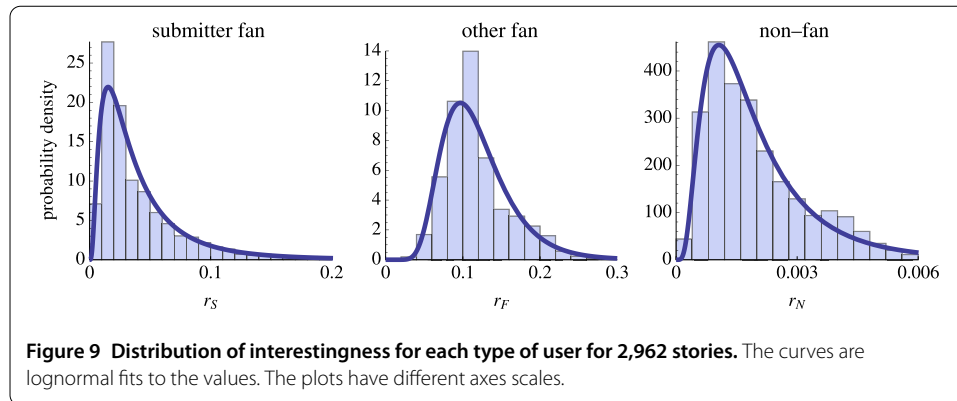


Table 2 Parameters for lognormal distribution of interestingness for 2,962 stories

Vote type	μ	σ
submitter fan	-3.48 ± 0.04	0.84 ± 0.03
other fan	-2.20 ± 0.02	0.37 ± 0.01
non-fan	-6.41 ± 0.03	0.66 ± 0.02

Error ranges are 95% confidence intervals for the parameter estimates.

to fans and non-fans, suggesting stories with particularly large ratios are mainly of niche interest.

Table 2 summarizes the lognormal distribution parameters. The lognormal fits provide a concise description of the distribution of interestingness values and are useful as prior distributions for prediction from early votes, as discussed in Section 5.2. However bootstrap tests [15] based on the Kolmogorov-Smirnov (KS) statistic show the estimated r -values are unlikely to arise from these distributions (p -values all less than 0.02). This test and the others reported in this paper account for the fact that we fit the distribution parameters to the data [16].

The relationship between interestingness for fans and other users indicates a considerable variation in how widely stories appeal to the general user community. Moreover, we find other fans have somewhat higher interest in stories than submitter fans, i.e., r_F tends to be larger than r_S . Since we have $c_S > c_F$ (Table 1), we find submitter fans are more likely to view the story while upcoming, but less likely to vote for it, compared with other fans. This suggests people favor the submitter as a source of stories to read, while the fact that a friend, not the submitter, voted for the story makes it more likely the user will vote for the story. Identifying these possibilities illustrates how models can suggest subgroups of behaviors in social media for future investigation.

5.2 Predicting popularity from early votes

In this section we use the model to predict popularity of Digg stories. We focus on the 89 of the 100 stories in the calibration data set that were promoted within 24 hours of submission. Most stories are promoted within 24 hours of submission (if they are ever promoted) and this restriction simplifies the model's use of the 'popular in last 24 hours' list by not requiring it to check for removal from the list if the story is still upcoming more than 24 hours after submission.

Predicting popularity in social media from intrinsic properties of newly submitted content is difficult [17]. However, users' early reactions provide some measure of predictability

[4, 11, 18, 19]. The early votes on a story allow estimating its interestingness to fans and other users, thereby predicting how the story will accumulate additional votes. These predictions are for expected values and cannot account for the large variation due, for example, to a subsequent vote by a highly connected user which leads to a much larger number of votes.

We can improve predictions from early votes by using the lognormal distributions of r -values, shown in Figure 9, as the prior probability to combine with the likelihood from the observations according to Bayes theorem. Specifically, instead of maximizing the likelihood of the observed votes, $P(r|\text{votes})$, as discussed above, this approach maximizes the posterior probability, which is proportional to $P(r|\text{votes})P_{\text{prior}}(r)$ where P_{prior} is taken to be the lognormal distribution $P_{\text{lognormal}}$ in Equation (14) with parameters from the fits shown in Figure 9.

For a prediction at time T , we use the votes up to time T to estimate the r values by finding the values that maximize

$$L = \log(P(r_S, r_F, r_N|\text{votes})) + \log(P_{\text{prior}}(r_S)P_{\text{prior}}(r_F)P_{\text{prior}}(r_N)). \quad (15)$$

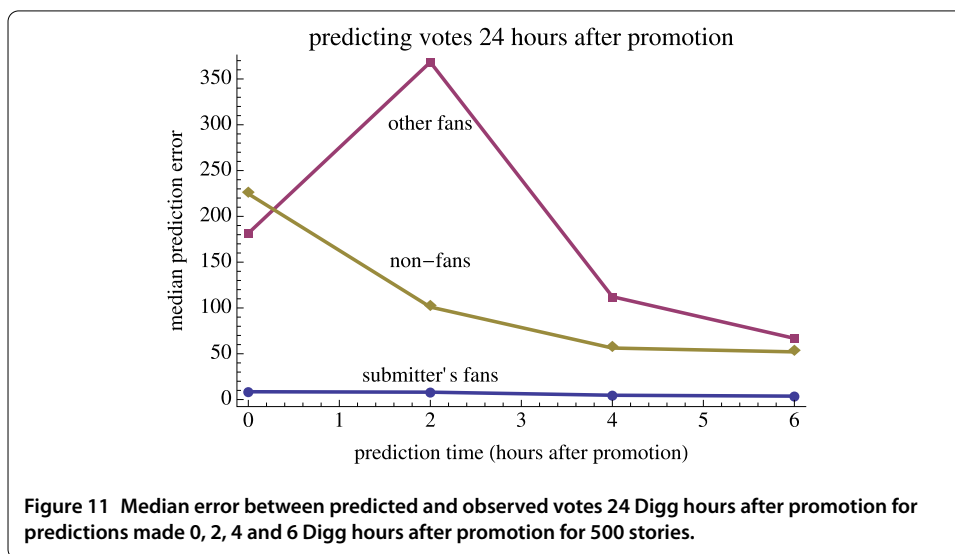
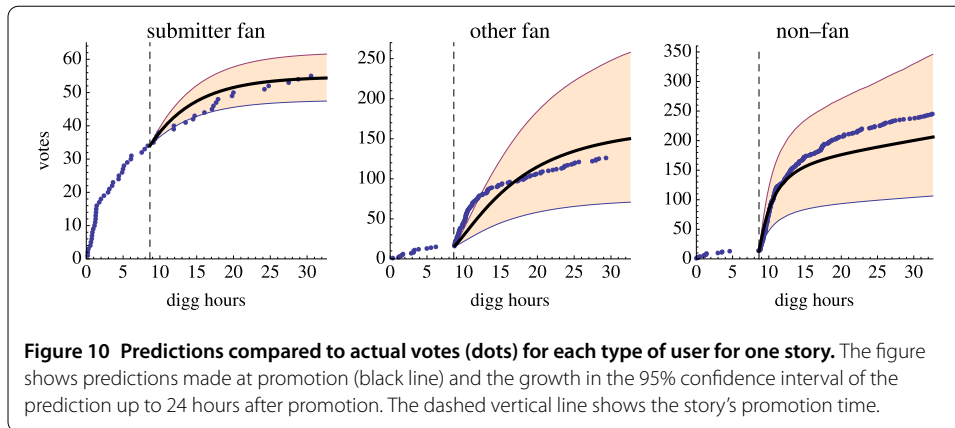
We then solve the model starting at time T and use the values from that solution as the predictions at later times. Solving the model equations starting at time T requires initial values, i.e., the number of votes $v_S(T)$, $v_F(T)$, $v_N(T)$ and the size of the user groups who have not yet seen the story: $S(T)$, $F(T)$, $N(T)$. The numbers of votes is available in our data but the sizes of the user groups is not. Instead, we estimate user group sizes from the voting rates and the estimated r values. That is, at the prediction time T we use the times of recent votes to estimate the derivative in number of each type of vote at time T . With estimates of the r values and voting rates, we use the model equations to estimate the group size at time T . For instance, Equation (1) gives

$$S(T) = \frac{1}{\omega r_S P_S} \frac{dv_S}{dt}. \quad (16)$$

We estimate the voting rates from the number of votes in the 15 minutes prior to time T , except if there are fewer than five votes in this time we extend the time interval to include the five previous votes. For simplicity, to avoid treating the discontinuity in visibility at promotion, we based this estimate on front page votes when T is after the promotion time. This approach to estimating user group sizes is reasonable while the story is accumulating votes rapidly, i.e., when it is recently promoted. In that case the story is highly visible and we have a good estimate of the voting rate. This is the situation of interest for prediction based on users' early response to a story. On the other hand, old stories have low visibility and accumulates votes slowly, if at all, as seen in Figure 1. In that case, the group size estimate, based on the ratio of voting rate and visibility, will be highly uncertain.

We focus on behavior after promotion. Figure 10 compares predicted to actual votes for one story 24 hours after promotion. Votes from submission to promotion are used to estimate r values for the three groups of users. The model solutions extend from the time of these estimates, i.e., the story's promotion time, to $t = 24$ Digg hours after promotion. The model quantitatively reproduces the observed votes for this story.

Generalizing from this example for a single story, Figure 11 shows the prediction errors for each type of vote the story receives 24 Digg hours after promotion, based on estimating



r -values from votes received by time T for various times T at and shortly after promotion. For context of the size of these errors, Figure 5 shows the range of number of votes the stories have at the times of these predictions, i.e., 24 hours after promotion.

As expected, errors generally decrease when predictions are made later. Of more interest is the difference among the type of votes, particularly for votes from other fans. Early votes are mainly from submitter's fans and non-fans, so the ability to predict differences in behavior for those groups based on early votes could be useful in quickly distinguishing stories likely to be of broad or niche interest to the user community.

Overall, the model reasonably predicts votes from submitter's fans and non-fans, but is much less accurate for votes from other fans. One reason for this difference is the relatively small number of other fan votes while a story is upcoming. Specifically, the number of other fans F starts at zero. Only a vote by a non-fan can increase F , and upcoming stories have low visibility to non-fan voters. Even after a number of other fans becomes available, it takes some time for those users to return to Digg. Thus there are relatively few early other fan votes, leading to poor estimates for r_F values. Moreover, the relatively small number of other fans means a single early voter with many fans can significantly change F away from its average value used in the model. These factors lead to the relatively large

Table 3 Spearman rank correlation between predicted and observed number of each type of votes 24 Digg hours after promotion, for predictions made at various times T after promotion (measured in Digg hours) for 500 stories

$T - T_{\text{promotion}}$	Correlation		
	Submitter fan	Other fan	Non-fan
0	0.86	0.29	0.42
2	0.94	0.74	0.91
4	0.97	0.78	0.93
6	0.97	0.81	0.95

Table 4 Classification errors on whether a story receives more than the median number of votes from each type of voter received by 24 Digg hours after promotion, for predictions made at various times T after promotion (measured in Digg hours) for 500 stories

$T - T_{\text{promotion}}$	Classification error		
	Submitter fan	Other fan	Non-fan
0	0.12	0.41	0.37
2	0.09	0.45	0.16
4	0.07	0.33	0.11
6	0.07	0.30	0.10

errors in predicting the other fan votes. As a direction for future work, this observation suggests predictions would benefit from including measurements of the social network of the voters to determine the value of F at the time of prediction rather than using an estimate based on the model.

Another view of prediction quality is how well the model predicts the rank ordering of stories, i.e., whether the story is likely to be relatively popular. We measure this with the Spearman rank correlation between the model's prediction and the observed number of votes 24 Digg hours after promotion, as shown in Table 3. Even for other fan votes, where the absolute prediction error is relatively large, the predicted values give a good indication of the relative rank of the stories.

Predicting whether a story will attract a large number of votes, rather than the precise number of votes, is a key issue for web sites such as Digg. Such predictions form the basis of using crowd sourcing to select a subset of submitted content to highlight [19]. As an example of this distinction, we predict whether a story will receive more than the median number of votes of each type of user based on votes received up to various times. This amounts to a binary classification task. Table 4 compares predictions made at different times. The classification error rate is the fraction of stories for which prediction of whether the story receives more than the median number of votes differs from the actual value.

5.3 Confidence intervals

We can use the model to estimate how well it predicts future votes. For a given set of parameter values, prediction variability comes from differences in estimated r values. If r is poorly determined, predictions will be unreliable.

To quantify this behavior, we numerically evaluate the second derivative matrix D of the log-likelihood combined with the priors based on votes on the story up to time T , $L(r)$ given in Equation (15), at the maximum $r = r_{\text{max}}$, where $r = (r_S, r_F, r_N)$. This gives

$$L(r) = L(r_{\text{max}}) + \frac{1}{2}(r - r_{\text{max}})D(r - r_{\text{max}}) \quad (17)$$

to second order in $|r - r_{\max}|$. To this order of expansion, the likelihood is

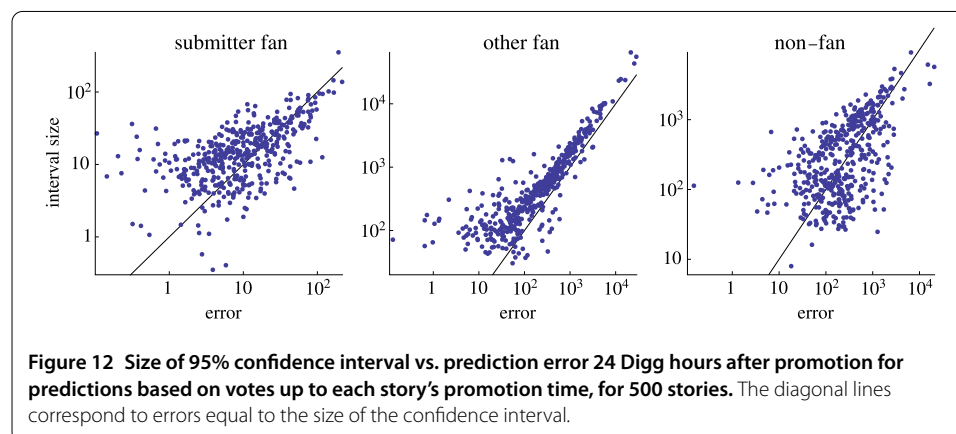
$$\exp(-(r - r_{\max})D(r - r_{\max})/2). \quad (18)$$

This corresponds to a multivariate normal distribution for r with mean r_{\max} and covariance matrix $-D^{-1}$. Since we are expanding around a maximum, the 2nd derivative matrix is negative definite so this gives a well-defined normal distribution, i.e., with a positive definite covariance matrix. This covariance includes both individual variances in the values of r_S , r_F and r_N and correlations among their variations around the maximum.

If $L(r)$ is a fairly flat function of r around the maximum, then maximum likelihood poorly constrains the values, corresponding to large variances in the normal distribution. Conversely, if $L(r)$ is sharply peaked, the distribution will be narrow.

We apply this observation to estimate confidence intervals for the predictions. We first numerically evaluate the second derivative matrix D at the maximum. We then generate random samples of r from the multivariate normal distribution. For each of these samples, we solve the model starting from the time T to any desired time for predicting the votes, e.g., 24 hours after promotion. After collecting these predictions from many samples, we use quantiles of their ranges as the confidence intervals. In the examples presented here, we generate 1,000 random samples and determine the 95% confidence interval from the variation in r values as the range between the 2.5% and 97.5% quantiles of these samples.

As one example, Figure 10 shows how confidence intervals grow with time after the prediction for predictions made from votes at the time a story is promoted. Predictions made at later times, when more votes are available, generally have smaller confidence intervals. More generally, Figure 12 shows the relation between 95% confidence interval and prediction error 24 Digg hours after promotion, based on prediction made at the time of promotion. Large errors tend to be associated with large confidence intervals. To quantify this relationship, correlations between prediction error and confidence interval size for the stories shown in Figure 12 are 0.76, 0.96 and 0.66 for submitter fans, other fans and non-fans, respectively. Thus the confidence intervals, which are computed from the vote information available at the time of prediction, indicate how well the model can predict votes. However, the confidence intervals based on variation in the estimate of r values do not account for all the sources of error. For instance, Figure 12 shows some cases where the error is considerably larger than the confidence interval. Furthermore, only about a third



of the actual votes 24 hours after promotion are within the confidence intervals, whereas we would expect about 95% of the stories to be within the intervals. Additional variation could be due, for instance, to votes by exceptionally well-connected users that significantly increase the story's visibility compared to the average value assumed with the model. Large prediction errors can also arise from poor estimates of the sizes of the groups of each type of user at the time of prediction.

6 Related work

Models of social dynamics can help explain and predict the popularity of online content. The broad distributions of popularity and user activity on many social media sites can arise from simple macroscopic dynamical rules [20]. A phenomenological model of the collective attention on Digg describes the distribution of final votes for promoted stories through a decay of interest in news articles [21]. Likewise, a model that accounted for shifts in public attention, e.g., brought about by exogenous events, reproduced in simulation the aggregate distribution of popularity of Web pages [22], but did not model dynamics of individual items. Yet another study [23] offered a qualitative explanation for the observed dynamics of popularity of individual news stories in terms of the dueling effects of preference for recent stories and those already featured in the news; however, no attempt was made to connect this model to empirical data. Stochastic models [1, 2] offer an alternative explanation for the popularity distribution of Digg news stories. Rather than novelty decay (or bias for recent news), they explain the votes distribution by the combination of variation in the stories' inherent interest to users and effects of user interface, specifically decay in visibility as the story moves to subsequent pages. Stochastic modeling framework explains both the dynamics of popularity of an individual news story, as well as the distribution of final popularity of many stories. Crane and Sornette [24] found that collective dynamics was linked to the inherent quality of videos on YouTube. From the number of votes received by videos over time, they could separate high quality videos from junk videos. This study is similar in spirit to our own in exploiting the link between observed popularity and content quality. We assume all users explore stories using the same parameters for the 'law of surfing' [8]. More generally, such models could be adjusted to accommodate a range of surfing persistence among different groups of users [25]. Web-based experiments allow directly manipulating visibility to distinguish its effect from social influence and content quality, e.g., by reversing the order in which content is shown to users [26]. State-based models, such as we used here, apply to forecasting user behavior in a wide variety of contexts, including web searches [27] which are one mechanism by which users could find content without relying on visibility explicitly provided by the social media site. However, while these studies aggregated data from thousands of individuals, our method focuses instead on the *microscopic* dynamics, modeling how individual behavior contributes to content popularity.

Statistically significant correlation between early and late popularity of content is found on Slashdot [18], Digg and YouTube [4]. Specifically, similar to our study, Szabo and Huberman [4] predicted long-term popularity of stories on Digg. Through large-scale statistical study of stories promoted to the front page, they were able to predict stories' popularity after 30 days based on its correlation with popularity one hour after promotion. Similarly, Lerman and Hogg [9] predicted popularity of stories based on their pre-promotion votes. We also quantitatively predict stories' future popularity, but unlike earlier works, we also estimate confidence intervals of these predictions for each story.

Previous works found social networks to be an important component to information diffusion. Niche interest content tends to spread mainly along social links in Second Life [28], in blogspace [29], as well as on Digg [19], and does not end up becoming very popular with the general audience. Models based on biased random walks to select actions [30] provide more detailed descriptions of information diffusion in social media [31] than used in our modeling approach. Aral et al. [32] found that social links between like-minded people, rather than causal influence, explained much of information diffusion observed on a network. Our modeling approach allows us to systematically distinguish users who are linked to those who are not linked and study diffusion separately for each group.

Commercial recommendation systems, such as those used by Amazon and Netflix, use collaborative filtering [33] to highlight new products for users. They ask users to rate products and compare ratings to identify users with similar opinions and recommend to them new products other users with similar opinions liked. Researchers have recognized that links between users in a recommender system can be induced from the declarations of user interest, and that these links can be used to make new recommendations [34]. However, unlike social media sites, collaborative filtering-based systems do not allow users to explicitly declare their social links. Nevertheless, such techniques can be useful in automatically helping social media users find others with similar interests [35–37] and thereby encourage users to make implicit links explicit in social media web sites. In our context, such methods could increase the visibility of similar users to each other, thereby improving the sorting of users between the fan and non-fan categories studied in our model.

7 Discussion

Highlighting friends' contributions is a common feature of social media sites, including Digg. To evaluate the effects of this behavior, we explicitly distinguish votes from submitter's fans, other fans and non-fans in our model, while separating the effects of differences in visibility and interestingness among these groups of users. This identifies that submitter's fans are, on average, far more likely to find the story interesting. Our model adjusts for the higher visibility of stories to fans, thereby identifying that increased attention from fans is not just due to the increased visibility. Identifying stories of particularly high interest to fans could be a useful guide for highlighting stories in the friends interface, i.e., emphasizing those with relatively large interestingness to friends as reflected in the early votes. Moreover, this information could be useful to recommend new fans to users, based on visibility-adjusted similarity in voting rather than, as commonly done in collaborative filtering [33], just using the raw score of similar votes. This could be particularly important for users with relatively infrequent votes, where variations due to how visible a story is could significantly affect the similarity of the vote pattern with that of other users.

For more precise estimates, the web site could track the fraction of users seeing the story that vote for it, thereby directly estimating interestingness and accounting for the large variability in number of fans among the voters, in contrast to our model which used an average value. Exploiting such details of user behavior becomes more important as the complexity of the web site interface increases, offering many ways for users to locate content. Recording which method leads each user to find the story can aid in identifying any systematic differences in interests among those users.

We find a wide range of interestingness ratios between fans and non-fans. This explains prior observations of the effect of relatively high votes from fans on indicating popularity

to the general user population, and also suggests stories that are of niche interest to the fans rather than the general user population. Our assumption that fans of prior voters easily see the story is reasonable for users with relatively few friends, so only a few stories will appear in their friends interface. For users with many friends, visibility of a story would decrease when many newer stories appear on the friends interface. This possibility could be included in the model using the 'law of surfing' for the stories appearing in each users' friends interface.

For prediction, we find the largest errors with votes from other fans. This likely arises from the relatively small number of such votes, especially while the story is upcoming. In that case, the large variation in number of fans per user can have a dramatic effect not accounted for in the model. This suggests the main source of the prediction error arises from the long-tail distribution of fans per user, which the model treats as a single average value based on the parameter ρ . We could test this possibility by collecting additional data on the actual fans of each voter, thereby using the observed value of $F(t)$ at the time of prediction when estimating r -values. In cases where $F(t)$ is particularly large, e.g., due to an early vote by a user with many fans, this will result in a smaller estimated value for r_F and hence smaller predicted number of other fan votes.

Models can suggest improved designs for user-contributory web sites. Our results suggest it may be useful to keep popular stories visible longer for users who return to Digg less often - giving them more chance to see the popular stories before they lose visibility. This would be a fine-tuned version of 'popular stories' pages, adjusted for each user's activity rate. That is, instead of showing stories in order of recency only, selectively move less popular stories down the page (once there are enough votes to determine popularity), thereby leaving the more popular ones nearer the top of the list for users who come back to Digg less often.

We examined behavior over a relatively short time (e.g., up to a day after promotion). Over longer times, additional factors could become significant, particularly a decrease in the interestingness as news stories submitted to Digg become 'old news' [21].

Modeling visibility depends on how the web site user interface exposes content. This highlights a challenge for modeling social media: continual changes to the user interface can alter how visibility changes for newly submitted content. Thus accurate models require not only data on user behavior but also sufficient details of the user interface at the time of the data to determine the relation between visibility and properties of the content.

The lognormal distribution of interestingness seen here and in other web sites [11] is useful as a prior distribution for estimating interestingness from early behavior on web sites. The use of such priors will be more important as models make finer distinctions among groups of users, e.g., distinguishing those who find the content in different ways as provided by more complex interfaces. In such cases, many groups will not be represented among the early reaction to new content and use of priors will be especially helpful.

User-contributory web sites typically allow users to designate others whose contributions they find interesting, and the sites highlight the activity of linked users. Thus our stochastic model, explicitly distinguishing behavior of users based on whether they are linked to users who submitted or previously rated the content, could apply to many such web sites.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

KL collected the data, TH performed the model parameter estimation. Both authors developed the model, wrote the paper, and approved the final manuscript.

Author details

¹Institute for Molecular Manufacturing, Palo Alto, CA, USA. ²USC Information Sciences Institute, Marina del Rey, CA, USA.

Acknowledgements

This work is supported in part by the Air Force Office for Scientific Research under contract FA9550-10-1-0102 and by the National Science Foundation under grant IIS-0968370. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies. Authors thank Suradej Intagorn for collecting and processing data.

Endnotes

- ^a At the time of data collection Digg offered a social filtering feature which recommended stories, including upcoming stories, that were liked by users with a similar voting history. It is not clear how often users employed these features and we do not explicitly include them in our model.
- ^b The data set is available at <http://www.isi.edu/~lerman/downloads/digg2009.html>.

Received: 7 February 2012 Accepted: 18 June 2012 Published: 18 June 2012

References

1. Lerman K (2007) Social information processing in social news aggregation. *IEEE Internet Comput* 11(6):16-28
2. Hogg T, Lerman K (2009) Stochastic models of user-contributory web sites. In: *Proc. of the third international conference on weblogs and social media (ICWSM2009)*. AAAI Press, Menlo Park, pp 50-57
3. Lerman K, Hogg T (2012) Using stochastic models to describe and predict social dynamics of web users. *ACM Transactions on Intelligent Systems and Technology* (to appear)
4. Szabo G, Huberman BA (2010) Predicting the popularity of online content. *Commun ACM* 53(8):80-88
5. Hogg T, Lerman K (2010) Social dynamics of Digg. In: *Proc. of the fourth international conference on weblogs and social media (ICWSM2010)*. AAAI Press, Menlo Park, pp 247-250
6. Haberman R (1987) *Mathematical models: mechanical vibrations, population dynamics, and traffic flow*. Classics in Applied Mathematics. SIAM, Philadelphia
7. Hethcote HW (2000) The mathematics of infectious diseases. *SIAM Rev* 42(4):599-653
8. Huberman BA, Pirolli PLT, Pitkow JE, Lukose RM (1998) Strong regularities in World Wide Web surfing. *Science* 280:95-97
9. Lerman K, Hogg T (2010) Using a model of social dynamics to predict popularity of news. In: *Proc. of the 19th intl. World Wide Web conference (WWW2010)*. ACM, New York, pp 621-630
10. Reed WJ, Jorgensen M (2004) The double Pareto-lognormal distribution: a new parametric model for size distributions. *Commun Stat, Theory Methods* 33:1733-1753
11. Hogg T, Szabo G (2009) Diversity of user activity and content quality in online communities. In: *Proc. of the third international conference on weblogs and social media (ICWSM2009)*. AAAI Press, Menlo Park, pp 58-65
12. Bulmer MG (1974) On fitting the Poisson lognormal distribution to species-abundance data. *Biometrics* 30:101-110
13. Miller G (2007) Statistical modelling of Poisson/log-normal data. *Radiat Prot Dosim*, 124:155-163
14. Hilbe JM (2008) *Negative binomial regression*. Cambridge Univ. Press, Cambridge
15. Efron B (1979) Bootstrap methods: another look at the jackknife. *Ann Stat* 7:1-26
16. Clauset A, Shalizi CR, Newman MEJ (2009) Power-law distributions in empirical data. *SIAM Rev* 51:661-703
17. Salganik M, Dodds P, Watts D (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311:854
18. Kaltenbrunner A, Gomez V, Lopez V (2007) Description and prediction of Slashdot activity. In: *Proc. 5th Latin American web congress (LA-WEB 2007)*
19. Lerman K, Galstyan A (2008) Analysis of social voting patterns on Digg. In: *Proceedings of the 1st ACM SIGCOMM workshop on online social networks*
20. Wilkinson DM (2008) Strong regularities in online peer production. In: *EC'08: Proceedings of the 9th ACM conference on electronic commerce*. ACM, New York, pp 302-309
21. Wu F, Huberman BA (2007) Novelty and collective attention. *Proc Natl Acad Sci USA* 104(45):17599-17601
22. Ratkiewicz J, Fortunato S, Flammini A, Menczer F, Vespignani A (2010) Characterizing and modeling the dynamics of online popularity. *Phys Rev Lett* 105(15):158701+. <http://prl.aps.org/pdf/PRL/v105/i15/e158701>
23. Leskovec J, Backstrom L, Kleinberg J (2009) Meme-tracking and the dynamics of the news cycle. In: *KDD'09: Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, New York, pp 497-506. <http://www.cs.cornell.edu/home/kleinber/kdd09-quotes.pdf>
24. Crane R, Sornette D (2008) Viral, quality, and junk videos on YouTube: separating content from noise in an information-rich environment. In: Lerman K et al (eds) *Proc. of the AAAI symposium on social information processing*, pp 18-20
25. Levene M, Borges J, Loizou G (2001) Zipf's law for web surfers. *Knowl Inf Syst* 3:120-129
26. Salganik MJ, Watts DJ (2008) Leading the herd astray: an experimental study of self-fulfilling prophecies in an artificial cultural market. *Soc Psychol Q* 71:338-355
27. Radinsky K, Svore K, Dumais S, Teevan J, Brocharov A, Horvitz E (2012) Modeling and predicting behavioral dynamics on the web. In: *Proc. of the 21st intl. World Wide Web conference (WWW2012)*
28. Bakshy E, Karrer B, Adamic LA (2009) Social influence and the diffusion of user-created content. In: *Proc. of the 10th ACM Conf. on electronic commerce (EC09)*, ACM, New York, pp 325-334
29. Colbaugh R, Glass K (2010) Early warning analysis for social diffusion events. In: *Proceedings of IEEE international conferences on intelligence and security informatics*

30. Bogacz R et al (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700-765
31. Götz M, Leskovec J, McGlohon M, Faloutsos C (2009) Modeling blog dynamics. In: Proc. of the third international conference on weblogs and social media (ICWSM2009). AAAI Press, Menlo Park, pp 26-33
32. Aral S, Muchnik L, Sundararajan A (2009) Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc Natl Acad Sci USA* 106(51):21544-21549. <http://dx.doi.org/10.1073/pnas.0908800106>
33. Konstan JA, Miller BN, Maltz D, Herlocker JL, Gordon LR, Riedl J (1997) GroupLens: applying collaborative filtering to Usenet news. *Commun ACM* 40(3):77-87. citeseer.ist.psu.edu/konstan97grouplens.html
34. Perugini S, Goncalves MA, Fox EA (2004) Recommender systems research: a connection-centric survey. *Journal of Intelligent Information Systems* 23(2):107-143
35. Backstrom L, Leskovec J (2011) Supervised random walks: predicting and recommending links in social networks. In: Proc. of the 4th ACM intl. conf. on web search and data mining (WSDM)
36. Bringmann B, Berlingerio M, Bonchi F, Gionis A (2010) Learning and predicting the evolution of social networks. *IEEE Intell Syst* 25(4):26-34
37. Schifanella R, Barrat A, Cattuto C, Markines B, Menczer F (2010) Folks in folksonomies: social link prediction from shared metadata. In: Proc. of the 3rd ACM intl. conf. on web search and data mining (WSDM2010)

doi:10.1140/epjds5

Cite this article as: Hogg and Lerman: Social dynamics of Digg. *EPJ Data Science* 2012 1:5.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com