J Stat Phys (2013) 151:329–339
DOI 10.1007/s10955-012-0596-5

Social Computing and the Attention Economy

Bernardo A. Huberman

Received: 31 August 2012 / Accepted: 12 September 2012 / Published online: 25 September 2012

Abstract Social computing focuses on the interaction between social behavior and information, especially on how the latter propagates across social networks and is consumed and transformed in the process. At the same time the ubiquity of information has left it devoid of much monetary value. The scarce, and therefore valuable, resource is now attention, and its allocation gives rise to an attention economy that determines how content is consumed and propagated. Since two major factors involved in getting attention are novelty and popularity, we analyze the role that both play in attracting attention to web content and how to prioritize them in order to maximize it. We also demonstrate that the relative performance of strategies based on prioritizing either popularity or novelty exhibit an abrupt change around a critical value of the novelty decay time, resembling a phase transition.

Keywords Social networks · Attention economics

1 Introduction

We are witnessing an inversion of the traditional way by which content has been generated and consumed over centuries. From photography to news and encyclopedic knowledge, the centuries old pattern has been one in which relatively few people and organizations produce content and most people consume it. With the advent of the web and the ease with which one can migrate content to it, that pattern has reversed, leading to a situation whereby millions create content in the form of blogs, news, ideas, videos, music, and product reviews, and relatively few can attend to it all. This phenomenon, which sometimes goes under the name of crowdsourcing, is exemplified by websites such as `digg.com`, Reddit, YouTube, Twitter, arXiv and Wikipedia, where content creation without the traditional quality filters manages to produce sought out movies, news and even knowledge that rival the best encyclopedias. That such content is valued is confirmed by the fact that access to these sites accounts for a sizable percentage of internet traffic. YouTube, which in many ways can be considered

B.A. Huberman (✉)
HP Laboratories, Palo Alto, CA 94304, USA
e-mail: bernardo.huberman@hp.com

a media company that outsources the production of its content to millions of users, has more viewers than all of the U.S. TV networks combined.

An interesting and daunting consequence of this inversion is that information, which used to be scarce and therefore valuable, is now so ubiquitous so as be almost devoid of monetary value. Search engines, billions of websites, targeted advertisement and easy access to digital content, provide us with myriad ways of taking care of our most complex informational and entertainment needs. What is now scarce, and therefore valuable, is the user's attention, which explains the intense efforts made at obtaining it through focused advertising, catchy titles of papers, pop-ups, artistic presentation of ideas in websites, and most disheartening, spam.

The attention economy, which focuses on how this scarce commodity is allocated among content and people, also plays a prominent role in the world of academia. Attention is often its main currency of the academic endeavour, for we publish to get the attention of others, we cite the work of colleagues so that they receive attention and we cherish the prominence of great work if only because of the attention it gathers [1, 2]. This is an old phenomenon that has been taking place since the establishment of learned societies and academic disciplines, but it is only recently that it is starting to be framed in the context of the new digital medium.

As millions of people use the web for their social, informational, and consumer needs, they also communicate with each other about their findings and interests, leading to the propagation of information through vast social networks that extend far beyond the traditional ones determined by geography and personal relations. At the same time, content providers vie for the limited attention of people by resorting to a number of strategies aimed at maximizing the amount of attention (usually measured in clicks) devoted to their web sites [5]. These strategies range from data personalization and short videos to the dynamic rearrangement of items in a given page, to name a few [6, 7, 12]. In all these cases the ultimate goal is the same: to draw the attention of the visitor of a site to its content before he proceeds to the next one [8].

Obviously, the more interesting and relevant the content of a site is, the more valuable it will be to users. In addition, since users need to decide among the existing plethora of links and sites, the popularity of given pieces of content are a determinant of their success, for people often click on given links for no other reason than the fact that many others do.

Another factor that determines success in the allocation of attention is novelty, for few users would keep attending to content that they viewed already or is stale in terms of its validity. That is why most news aggregators, and sites ranging from YouTube to arXiv order the links of a given page by their novelty, so as to guarantee a high degree of attention.

Within this context, we have recently shown that there is indeed a strong interplay between novelty and collective attention, which is universally manifested in a rather swift initial growth of the number of people looking at a new item within a site and its eventual slowdown as interest fades among the population [11].

And yet, given the role that popularity also plays in attracting the attention of users, a natural question within the attention economy is whether alternative orderings, like one giving priority to popularity over novelty, might not do better at attracting viewers to a site and thus maximizing attention.

We answered this question by taking the dynamics of collective attention to a finer level of detail and examined the role that popularity and novelty play in determining the number of clicks within a given page [3]. Specifically, we studied three different strategies that can be deployed in order to maximize attention. The first strategy prioritizes novelty while the second emphasizes popularity. The third strategy looks myopically into the future and prioritizes stories that are expected to generate the most clicks in the next few minutes. We show that the first two strategies should be selected on the basis of the rate of novelty decay, while the third strategy performs sub-optimally in most cases. Most interestingly from a statistical physics point of view, we discovered that the relative performance of the first two benchmark strategies as a function of the rate of novelty decay switches so sharply around some critical value so as to resemble a phase transition.

This paper is organized as follows. We first consider the question of whether or not the location of a link in a page determines the overall number of clicks in a given time interval. Having answered this in the affirmative through an empirical study of `digg.com`, we then proceed to introduce a set of indexes whose values determine the optimal strategy to be pursued in order to maximize attention to a page. Using the measured values of the rate of decay from `digg.com` we built a realistic simulator to collect statistically significant data to measure the performance of each of the indices introduced.

We then study the performance of each of these indices as a function of the novelty decay rate and show which strategy optimizes viewing for given values of the decay. Most importantly, we compute a full phase diagram that indicates at a glance the optimal strategy to use given the parameter values of the site. This phase diagram exhibits a sharp boundary between the choice of prioritizing novelty over popularity, thus resembling a phase transition.

Finally we summarize our results and discuss their implications for the design of dynamic websites.

2 Location Matters

In this section we study how the order in which links are placed within a webpage (e.g. the news stories of `digg.com`) determines the number of clicks within a certain time frame. Assume that time flows discretely as $t = 0, 1, 2, \dots$ minutes. Let N_t denote the number of clicks, or *digg number* of a story in `digg.com`, that appeared on the website t minutes ago (in this case we say that the story has *lifetime* t). As we showed earlier [11] the growth of N_t through transmission over a social network satisfies the following stochastic equation:

$$N_{t+1} = N_t(1 + a r_t X_t), \tag{1}$$

where r_t is a *novelty factor* that decays with time and satisfies $r_0 = 1$, X_t is a random variable with mean 1, and a is a positive constant.

This equation takes into account two important factors that together determine the growth of collective attention: *popularity* and *novelty*. The popularity effect is captured by the multiplicative form of Eq. (1), and the novelty effect is described by r_t. All other factors are contained in the noise term X_t.

The asymptotic behavior of this stochastic equation predicts a lognormal probability distribution for N, which been confirmed by the analysis of large data sets from `digg.com` and YouTube [4, 11]. Moreover its decay follows a nearly exponential function, suggesting a natural time scale for the decay of attention.

We next take the analysis to a finer level by considering a third *position factor*. A news story displayed at a top position on the front page easily draws more attention than a similar story placed on later pages. Hence the growth decay $a r_t$ should depend on the physical position at which the story is posted.

In the specific case of `digg.com`, its front page is divided into 15 slots, being able to display 15 stories at a time. The stories are always sorted chronologically, with the latest story at the top. If we label the positions from top to bottom by $i = 1, 2, \dots, 15$, we can modify Eq. (1) to allow for an explicit dependency of a on i:

$$N_{t+1} = N_t(1 + a_i r_t X_t), \tag{2}$$

where a_i is a position factor that decreases with i.

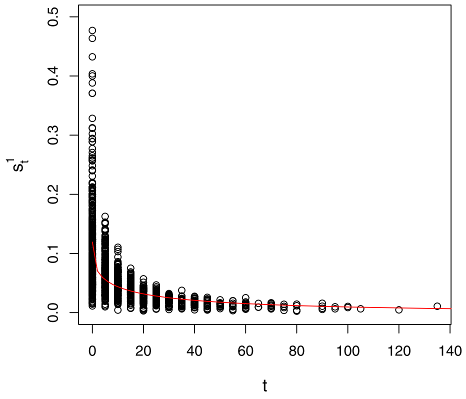

(a)

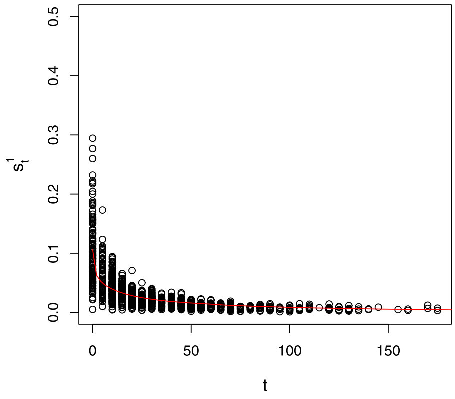

(b)

Fig. 1 The logarithmic growth rate for the top two positions on the front page of `digg.com`. Time is measured in minutes. Data is collected every 5 minutes, the rate at which the front page is refreshed. The *solid curve* in (**a**) is the result of a minimum mean square fit to the data (see text for more details). It has the functional form $f(t) = 0.120\,e^{-0.4t^{0.4}}$. The curve in (**b**) has the functional form $f(t) = 0.106\,e^{-0.4t^{0.4}}$

The assumption that the novelty effect and the position effect can be separated into two factors r_t and a_i needs to be tested empirically. To this end we tracked the growth rate for each slot, rather than for each story. For multiplicative models it is convenient to define the logarithmic growth rate

$$s_t = \log N_{t+1} - \log N_t. \tag{3}$$

When a is small (which is always true for short time periods) we have from Eq. (2)

$$s_t^i \approx a_i r_t X_t \tag{4}$$

for a story placed at position i at time t. Taking expectation of both sides, we have

$$Es_t^i \approx a_i r_t, \tag{5}$$

since $EX_t = 1$.

The logarithmic growth rate s_t^i can be measured as follows. For each fixed position i, if a `digg` story appears on that position at both times t and $t+5$ (the front page is refreshed every 5 minutes), then the observed quantity $\frac{1}{5}(\log N_{t+5} - \log N_t)$ counts as one sample point of s_t^i. Figure 1(a) plots 1,220 sample points collected from the top position at various times. Figure 1(b) is a similar plot for the second top position. By comparing (a) and (b) we see that s_t^2 indeed tends to fall below s_t^1, which indicates that the position effect is real. To better illustrate the position effect, we plot the expected growth rate for position 1, 3 and 5 in Fig. 2. As can be seen there, the growth rate decays as the story moves to lower positions.

From this data we can also determine the values of a_i quantitatively. We already established that for `digg.com` the precise functional form of the decay factor is $r_t = e^{-0.4t^{0.4}}$. Thus, for these particular values, the minimum mean square estimator $\hat{a}^i$ minimizes

$$\min_{a^i} \sum_j \left[s_{t_j}^i(j) - a^i r_{t_j}\right]^2 = \min_{a^i} \sum_j \left[s_{t_j}^i(j) - a^i e^{-0.4t_j^{0.4}}\right]^2, \tag{6}$$

where t_j is the lifetime of the j'th data point. The estimator for the 1,220 data points obtained from the top position is calculated to be $\hat{a}^1 = 0.120$. The fitted curve

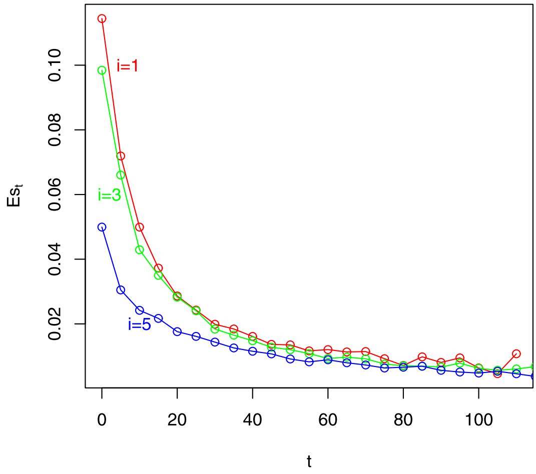

Fig. 2 The expected logarithmic growth rate for position 1, 3 and 5 on the front page of `digg.com`. Time is measured in minutes. As can be seen, the growth rate decays as the story moves to lower positions

$\hat{a}^1 r_t = 0.120 e^{-0.4 t_j^{0.4}}$ is shown as a solid curve in Fig. 1(a). An estimator $\hat{a}^2 = 0.106$ for the second top position is also calculated and plotted in Fig. 1(b). As can be seen from those figures, the position effect (a^i) and the novelty effect (r_t) can indeed be separated. We can then conclude that Eq. (2) fits the data very well.

3 Optimal Ordering for Maximal Attention

We now consider the order in which news stories should be displayed on a web page so as to generate the largest number of clicks within a certain time period T. This time period needs to be finite because the total number of clicks diverges as T goes to infinity. Equivalently, in an infinite-horizon framework, we could discount future clicks with a discount parameter δ, so that one click at time t counts as δ^t click at time 0. The objective then is to maximize $\sum_{t=0}^{\infty} \delta^t N_t$, where N_t is the total number of clicks generated from the news page in period t. In what follows we will consider the finite-horizon objective.

To simplify the problem we confine ourselves to a subset of ordering strategies called *indexing strategies*, which is defined as follows. Given a story's state, which in our model is just a two-vector (N_t, t), one first calculates an index O for each story using a predefined *index function* $O(N_t, t)$, and then sorts the stories based on their indices. The story with the largest index is displayed at the top, the story with the second largest index next, and so on [9, 10].

Rather than considering a general index function we will concentrate on three simple strategies. While neither of them is perfect, each can increase overall attention to the site.

1. $O_1(t) = -t$. The stories are sorted by their novelty, with the newest story at the top. This is what `digg.com` is doing today.
2. $O_2(t) = N_t$. The stories are sorted by their popularity, with the most popular story at the top. This strategy is based on the fact that attention grows in a multiplicative fashion (popular stories are more likely to become even more popular).
3. $O_3(t) = N_t r_t$. This is the "one-step-greedy" strategy. Ignoring the position effect (assume $a = 1$), a story in state (N_t, t) generates on average $N_t r_t$ more clicks (or "diggs" if one considers `digg.com`) in the next period. This strategy thus places the most "replicated" story at the top.

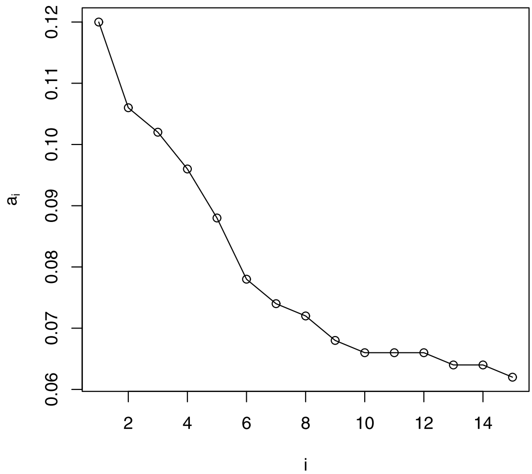

Fig. 3 The position factor decays as the position lowers. The values of a_i are measured by tracking the 15 slots on `digg.com`'s front page

Notice that because N_t grows with time, the effect of sorting by O_1 is almost the opposite of sorting according to O_2.

In order to test these strategies, we built a simulator that closely resembles the functioning of `digg.com` in that it incorporates the following rules:

1. Initially there are 15 stories, all in state $(N_t, t) = (1, 0)$. In words, each story starts with 1 digg and lifetime 0. (Because our model is purely multiplicative, the initial digg number does not matter. We just set it to be 1.)
2. Allocate the 15 stories to 15 positions, in decreasing order of their $O(N_t, t)$, for any given index function O.
3. Time evolves one step (5 minutes) at a time. The number of diggs generated from a story at position i is given by
$$\Delta N_{t+5} = N_{t+5} - N_t = 5 a_i r_t X_t N_t. \tag{7}$$
The total number of diggs generated in this time step is the sum of 15 such numbers. The values of a_i were estimated from real data and shown in Fig. 3. $r_t = e^{-0.4t^{0.4}}$. X_t is randomly drawn from a normal distribution with mean 1 and standard deviation 0.5 (obtained from the real data from `digg.com`).
4. On average every 20 minutes a new story arrives. Thus the number of stories arriving in one time step (5 minutes) follows a Poisson distribution with mean 0.25. When a new story enters the pool, the story with the lowest index is dropped, maintaining 15 stories in total. (It is possible the a new story is dropped immediately after its arrival if it happens to have the lowest index.)
5. Go back to Step 2 until the loop has been repeated for enough rounds.

The performance of all three index functions were tested in our simulator. For each index function, Steps 2 to 5 were repeated 100,000 times (or equivalently 500,000 minutes). Strategy O_1 (sort by novelty) achieved a total number of 514,314.8 diggs. Strategy O_2 (sort by popularity) only generated 354.6 diggs. Strategy O_3 (one-step-greedy) generated 452,402.3 diggs. Thus for these parameter values O_1 turns out to be best strategy, since it is 13.7 % better than O_3 and tremendously better than O_2. This confirms that `digg.com` is using the right strategy.

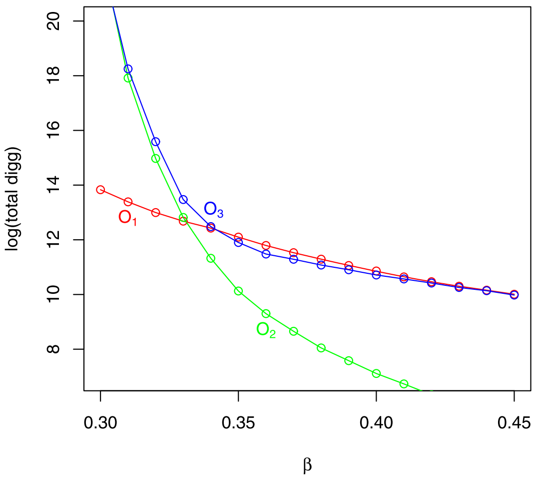

Fig. 4 The total number of diggs generated using three ordering strategies O_1, O_2, and O_3, for $\alpha = 0.4$ and a range of β. The novelty factor decays as $r_t = e^{-\alpha t^\beta}$. Performance is measured by the logarithm of the total number of diggs generated in 10,000 time steps. As can be seen, O_3 asymptotically approaches O_1 and O_2 in the fast decay (large β) and slow decay (small β) limit, respectively. A phase transition happens around $\beta = 0.335$

The reason for the poor performance of the index O_2 is easy to understand. O_2 gives higher priority to stories that have been dugg many times. According to the indexing rule, after one period new stories can never find their way to the front page since all the old stories have more than 1 digg! When novelty decays fast, the old stories remaining on the front page soon lose their freshness and cease to generate any new diggs. The system thus gets frozen in an unfruitful state.

The fact that O_1 outperforms O_3 is a bit harder to understand. Some intuition can be gained by considering an extreme case. Suppose each story completely loses its novelty after one second ($r_0 = 1$, $r_t = 0$ for all $t > 0$). Then only "new arrivals" should be displayed since they are the only ones that can generate new diggs. Sorting stories by their lifetime is a good idea when novelty decays fast. On the other hand, if novelty never decays ($r_t \equiv 1$), the lifetime factor becomes irrelevant. Thus in this case, strategy O_3, which prioritizes popular stories, will win over O_1. Hence, the fact that O_1 works better than O_3 in our simulations shows that novelty decays relatively fast for `digg.com`. Should it decay at a slower rate, O_3 would be a better choice.

We point out that our simulation only showed that the ordering implied by O_1 works better than O_3 for a particular choice of T. In general this may not be true for other values of T. In fact, for a time interval of $T = 5$ minutes (one time step) O_3 is by definition the best strategy. Hence, comparing the performance of two or more index functions only makes sense after one has specified a time horizon (or how much the future should be discounted if an infinite horizon is assumed).

In order to quantitatively test the limiting behavior of the three strategies, we repeated our simulations for a range of different values of the decay parameter r_t. Our previous work suggested that r_t decays as a stretched exponential function, whose general form can be written as $r_t = e^{-\alpha t^\beta}$. For `digg.com` it turns out $\alpha = \beta = 0.4$. The parameter β determines the decay rate. For fixed α, the larger β, the faster r_t decays. We repeated our experiment for $\alpha = 0.4$ and $\beta \in [0.30, 0.45]$. The result is shown in Fig. 4. The performance of each indexing strategy is measured by the logarithm of the total number of diggs generated in 10,000 rounds. We see that as β increases (faster decay), the number of diggs decreases for all three indexing strategies. When $\beta > 0.34$, O_1 performs slightly better than O_3 and much better than O_2. When $\beta < 0.33$, however, O_3 and O_2 perform significantly better than O_1. In other words, on the two sides of the value of $\beta = 0.335$, the stories should be displayed in completely reversed order! We therefore say that a *phase transition* takes place at the value of $\beta = 0.335$.

Other points worth mentioning are that in Fig. 4 O_3 asymptotically approaches O_1 and O_2 both in the fast and slow decay limits, and that in general O_3 is the best index among the three strategies (although for the specific parameters of `digg.com` ($\alpha = \beta = 0.4$) and our particular time horizon O_1 is slightly better). This is because O_3 trades off between popularity and novelty instead of betting on only one factor. To see this, consider the equivalent index function

$$O'_3(N_t, t) = \log O_3(N_t, t) = \log N_t + \log r_t. \tag{8}$$

Clearly, O'_3 linearly trades off between $\log N_t$ and $\log r_t$, assigning identical weight to the two effects. This is by no means the best tradeoff. For example, the index function

$$O_4(N_t, t) = 0.6 \log N_t + \log r_t \tag{9}$$

achieves 556,444.1 diggs after 100,000 rounds of simulation, which is 8.2 % more than O_1 and 23.0 % more than O_3! However arbitrary it may seem to give the term $\log N_t$ weight 0.6 rather than 1 is beyond the scope of this paper, but it does show the complexity of our problem. These experiments demonstrate that the novelty decay rate needs to be measured with great care, as a slight change in the decay rate may totally reverse the optimal order needed to maximize attention.

It is usually hard to analytically compute the performance of a general index function. For the two simple strategies O_1 and O_2, however, some rough estimate can be achieved. For the sake of generality, assume that there are m positions on the front page. New stories arrive at a rate $\lambda > 0$. Novelty decays as $r_t = e^{-\alpha t^\beta}$, where $0 < \beta \le 1$. Let $\bar{a} = \frac{1}{m}\sum a_i$ be the average position factor, which equals 0.08 for `digg.com`. Let Δt be the refresh time step, which is 5 minutes for `digg.com`.

Consider strategy O_2 first. According to the index rule, new stories never appear on the front page. All diggs are generated by the initial m stories. After time T we have from Eq. (3) that

$$\log N_T = \sum_{t=0,\Delta t,\ldots,T-\Delta t} a_i r_t X_t \Delta t. \tag{10}$$

Hence on average each story's log-performance is

$$E \log N_T = \sum_{t=0,\Delta t,\ldots,T-\Delta t} \bar{a} r_t \Delta t \approx \bar{a} \int_0^T r_t dt. \tag{11}$$

When T is large, we have

$$E \log N_T \approx E \log N_\infty = \bar{a} \int_0^\infty r_t dt. \tag{12}$$

Next consider O_1, which orders the stories by their lifetime. On average every $s \equiv 1/\lambda$ minutes a new story replaces an old story, and each old story moves down one position. Hence on average each story stays on the front page for ms minutes, where m is the number of positions. We call ms one *page cycle*. It is the average time it takes to refresh the whole page. We now see that, before a story disappears from the front page, it generates

$$N_{ms} = \exp\left(\sum_{t=0,\Delta t,\ldots,ms-\Delta t} a_{i(t)} r_t X_t \Delta t \right) \tag{13}$$

diggs, where $i(t)$ is the story's position at time t. When an story gets replaced by a new story, they are counted as one story restarting from the state $N_t = 1$ and $t = 0$. The multiplicative process starts over, and another N_{ms} diggs are generated in the next ms minutes, on average. Thus, in a total time period T the process is repeated $T/(ms)$ times, and a total number of $N_{ms}T/(ms)$ diggs are generated per story. The log-performance of O_1 is approximately

$$\log N_{ms} + \log\left(\frac{T}{ms}\right) = \sum_{t=0,\Delta t,\ldots,ms-\Delta t} \bar{a} r_t X_t \Delta t + \log\left(\frac{T}{ms}\right), \tag{14}$$

where we replaced $a_i(t)$ by $\bar{a}$ since on average each story stays in position $1, \ldots, m$ for equal times. Taking expectation on both sides, we have

$$E \log N_{ms} + \log\left(\frac{T}{ms}\right) \approx \bar{a} \int_0^{ms} r_t dt + \log\left(\frac{T}{ms}\right). \tag{15}$$

The critical point can be determined by equating Eqs. (12) and (15):

$$E \log N_T - E \log N_{ms} = \log T - \log(ms), \tag{16}$$

or

$$\bar{a} \int_{ms}^{\infty} r_t dt = \log\left(\frac{T}{ms}\right), \tag{17}$$

which holds for any functional form of r_t. The left side of Eq. (16) can be interpreted as the total novelty left after a time ms, or the total log-performance that can be gained from one story after one page cycle. The right hand side of Eq. (16) is the total log-time left after one page cycle. Thus, Eqs. (16) and (17) say that, after one page cycle, if there is more novelty left than the log-time remained, the stories should be ordered by decreasing popularity rather than by decreasing novelty (O_2 is better than O_1). Conversely, if novelty decays too fast (not enough novelty left after one page cycle), then the stories should be ordered by decreasing novelty rather than decreasing popularity (O_1 is better than O_2).

When $r_t = e^{-\alpha t^\beta}$ it holds that

$$\int_{ms}^{\infty} r_t dt = \frac{\alpha^{-\frac{1}{\beta}}}{\beta} \Gamma\left(\frac{1}{\beta}, \alpha(ms)^\beta\right), \tag{18}$$

where

$$\Gamma(a, x) = \int_x^{\infty} t^{a-1} e^{-t} dt \tag{19}$$

is the *incomplete Gamma function*. In this case the critical equation can also be written as

$$\bar{a} \frac{\alpha^{-\frac{1}{\beta}}}{\beta} \Gamma\left(\frac{1}{\beta}, \alpha(ms)^\beta\right) = \log\left(\frac{T}{ms}\right). \tag{20}$$

For the parameters of `digg.com` ($\bar{a} = 0.08$, $m = 15$, $s = 20$) and horizon $T = 50{,}000$ one can solve for the critical curve (α, β) on which O_1 and O_2 have the same performance. The curve is shown in Fig. 5 as a phase diagram. When the parameters (α, β) lie above the critical curve, the stories should be sorted by O_1. Otherwise they should be sorted by O_2.

To illustrate how sharp the phase transition is, we plot the relative performance $O_2/(O_1 + O_2)$ as a function of β, for fixed $\alpha = 0.4$, in Fig. 6. As can be seen, the transition is indeed very sharp.

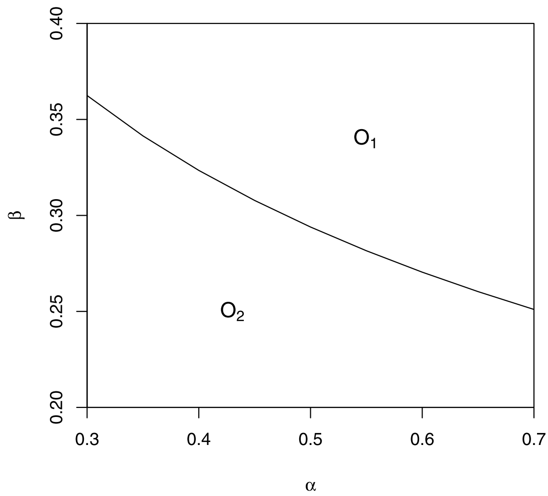

Fig. 5 The phase diagram. The critical curve is calculated by solving Eq. (20) with $\bar{a} = 0.08$, $m = 15$, $s = 20$ and $T = 50{,}000$. When (α, β) lies in the upper half of the phase diagram O_1 works better than O_2. Otherwise O_2 works better

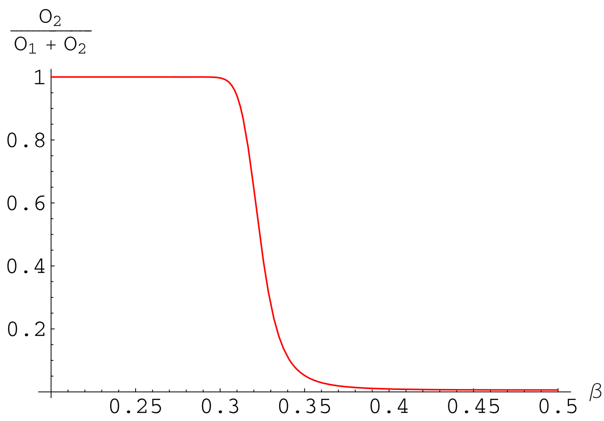

Fig. 6 The relative performance $O_2/(O_1 + O_2)$ as a function of β, for fixed $\alpha = 0.4$

4 Conclusion

In the attention economy, a top priority for those who create content is to successfully compete for the attention of others. Given the skewed fashion by which attention is allocated it is by no means trivial to design strategies to accomplish this goal. But given that within social media popularity and attention play such an important role at attracting attention, it is natural to ask how to prioritize them so as maximize the visibility of given content.

As we showed, depending on the rate of decay of novelty, two different strategies can be deployed in order to maximize attention. The first one prioritizes novelty while the second emphasizes popularity. Most interestingly, the shift from one to the other as a function of the rate of decay is extremely sharp, resembling a phase transitions.

These results were obtained by focusing on the dynamics of collective attention and examining the role that popularity and novelty play in determining the number of clicks within a given page. In particular, we analyzed three different strategies that can be deployed in order to maximize attention. The first strategy prioritizes novelty while the second emphasizes popularity. The third strategy looks myopically into the future and prioritizes stories that are expected to generate the most clicks in the next few minutes. We then showed that the first two strategies should be selected on the basis of the rate of novelty decay, while the third strategy performs sub-optimally in most cases. Most interestingly, we discovered that the relative performance of the first two benchmark strategies as a function of the rate of novelty decay switches so sharply around some critical value that it resembles phase transitions observed in the real world.

Given the widespread role that the attention economy plays in the success and failure of ideas, products and trends, these methods provide ways of maximizing the value that both content providers and users can obtain in information rich environments which vastly exceed our capacity to search and process on our own.

References

1. Franck, G.: Science communication, a vanity fair. Science **286**, 53–55 (1999)
2. Klamer, A., Dalen, H.P.V.: Attention and the art of scientific publishing. J. Econ. Methodol. **9**, 289–315 (2002)
3. Wu, F., Huberman, B.A.: Popularity, novelty and attention. In: Proceedings of the 2008 ACM Conference on Electronic Commerce, 2008
4. Szabo, G., Huberman, B.A.: Predicting the popularity of online content. Commun. ACM **8**, 80–88 (2010)
5. Falkinger, J.: Attention economies. J. Econ. Theory **133**, 266–294 (2007)
6. Garofalakis, J., Kappos, P., Mourloukos, D.: Web site optimization using page popularity. IEEE Internet Comput. **3**(4), 22–29 (1999)
7. Hong, W., Thong, J.Y.L., Tam, K.Y: Does animation attract online users attention? The effects of flash on information search performance and perceptions. Inf. Syst. Res. **15**(1), 60–86 (2004)
8. Huberman, B.A., Pirolli, P.L.T., Pitkow, J.E., Lukose, R.M.: Strong regularities in world wide web surfing. Science **280**(5360), 95–97 (1998)
9. Niño-Mora, J.: Stochastic scheduling. In: Floudas, C.A., Pardalos, P.M. (eds.) Encyclopedia of Optimization, vol. V, pp. 367–372 (2001)
10. Wu, F., Huberman, B.A.: The economics of attention: maximizing user value in information-rich environments. In: The First International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD'07), 2007
11. Wu, F., Huberman, B.A.: Novelty and collective attention. Proc. Natl. Acad. Sci. **104**(45), 17599–17601 (2007)
12. Zhang, P.: The effects of animation on information seeking performance on the world wide web: securing attention or interfering with primary tasks? AIS **1**(1) (2000)