I. Troch, F. Breitenecker, eds. ISBN 978-3-901608-35-3

# AN INDIVIDUAL-BASED MODEL OF COLLECTIVE ATTENTION

M.Moussaid[1 2], D. Helbing[2], G. Theraulaz[1]

[1]Centre de Recherches sur la Cognition Animale, CNRS, Université Paul Sabatier, Toulouse, France; [2]ETH Zurich, Swiss Federal Institute of Technology, Chair of Sociology, Zürich, Switzerland.

Corresponding Author: M.Moussaid, Centre de Recherches sur la Cognition Animale, UMR 5169, CNRS – Université Paul Sabatier, Bât. 4R3, 118 route de Narbonne, 31062 Toulouse cedex 09, France; mehdi.moussaid@gmail.com

**Abstract**. In our modern society, people are daily confronted with an increasing amount of information of any kind. As a consequence, the attention capacities and processing abilities of individuals often saturate. People, therefore, have to select which elements of their environment they are focusing on, and which are ignored. Moreover, recent work shows that individuals are naturally attracted by what other people are interested in. This imitative behaviour gives rise to various herding phenomena, such as the spread of ideas or the outbreak of commercial trends, turning the understanding of collective attention an important issue of our society.
In this article, we propose an individual-based model of collective attention, considering two recently measured empirical laws: First, an observation by Stanley Milgram describing how the interest of some people stimulates the interest of others, and, second, recent findings regarding how the interest in a news item decays over time.
In a situation where a group of people is facing a steady flow of novel information, the model naturally reproduces the log-normal distribution of the attention each news item receives, in agreement with empirical observations of the dynamics in place on the website *digg.com*. Furthermore, the model predicts that the popularity of a news item strongly depends on the number of concurrent news appearing at approximately the same moment. We confirmed this prediction by means of empirical data extracted from the website *digg.com*. This result can be interpreted from the point of view of competition between the news for the limited attention capacity of individuals. The proposed model, therefore, provides new elements to better understand the dynamics of collective attention in an information-rich world.

## 1 Introduction

With the ongoing growth of mass media and communication technologies, people are every day confronted with an increasing amount of information of any kind. Pieces of information can be broadcasted through television, urban advertisements or Internet, or exchanged directly between people through emails, phone calls or personal communications.
As a side effect, the amount of information an individual is facing in its everyday life exceeds by far its processing abilities and attention capacities [1]. Therefore, people have to *select* which elements of their environment they are focusing on, and which are ignored or checked out roughly. Interestingly, when selecting items they pay attention to, individuals are strongly influenced by other people's choices. First, because people like to share the same topics of interest as their friends, neighbors or colleagues, second, because popular novels are more relayed in the media, which increases their level of attraction and, finally, because people are naturally curious about what others are interested in [2].
The study of collective attention, therefore, is an important issue in the understanding of various herding phenomena, such as the spread of ideas [3], the dynamics of donations [4], or the outbreak of commercial trends resulting in bestsellers or blockbusters [5].
Recent work shows that some patterns of collective attention can be understood by invoquing epidemic models. For example, the temporal evolution of video views on *YouTube.com* is well described by a macroscopic model, where first viewers trigger a cascade of subsequent views over a social network [5]. Similarly, the donation rate after the Asian tsunami and various aspects of collective attention on the interactive website *digg.com* can be understood by means of similar processes [4][6].
However, the literature generally lacks individual-based models of collective attention. In contrast to the macroscopic description, individual based models start from a description of local interactions between people. Numerical simulations of the model should then produce the emergent properties of the system under study. Such an approach is valuable, because it allows the exploration of the phenomena from a different perspective, and complements well the macroscopic approaches previously adopted.

In this contribution, we therefore propose an individual-based model of collective attention, describing how people behave and interact when facing a steady flow on novel information. The paper is organized as follow: We first describe in more details the phenomena under study and highlight some empirical patterns of collective attention, as observed on the World Wide Web. Then follows a description of the model that we suggest to explain the above patterns. Finally, we present first simulation results and see how well they match the above observations and what we can learn from them.

## 2 Patterns of collective attention

In the age of the World Wide Web, scientists have gained access to unprecedented volumes of data on human activities [5-8]. Evaluating people's activity on the Internet constitutes a relevant way to observe patterns of collective attention as well. The interactive website *digg.com*, for example, is an interesting source of data to observe peoples' behaviors in situations where they have to select relevant information from a very large pool of items. The website allows its users to submit contents found elsewhere on the web. Each submitted story appears on the website and is dynamically assigned to an explicit measure of popularity, which motivates people to read it or not. If a user likes a story, he can add a "digg" to it. The total number of diggs a story has received is displayed next to it, and provides an indicator of its level of popularity[1]. If a submission receives enough diggs within a certain time period, it eventually jumps to the 'popular' section and becomes accessible by the users more easily. In practice, a large number of users explore and interact within this popular section of the website, which gives rise to an interesting collective dynamics.

The website offers a free access to the database of stories. It is therefore possible to track the temporal dynamics of each item's popularity to see how it has attracted the attention of readers over time [6]. To investigate this dynamics, we tracked the evolution of the digg rate for 1137 popular stories submitted from July 2 to July 12, 2007. The digg rate $\gamma$ is defined as the number of diggs that a story has received during a certain period of time $\delta t$ (here we choose $\delta t$=10 minutes). This enabled us to track the evolution of a story's popularity over time (Fig. 1a). The average curve displays a typical popularity pattern, characterized by a sudden burst of attention shortly after the story comes out, followed by a slow relaxation to low popularity (Figure 1a). Alternatively, one can observe how often a specific keyword has been searched on *Google* over time[2] [5]. Although very relevant, this latter method remains limited because one has to find enough keywords to draw a statistically consistent average pattern. As an illustration, the search pattern for the keyword '*tsunami'* during the Asian catastrophe in 2004 is displayed in the inset of figure 1a. We observe a very similar dynamics, suggesting that this pattern is not specific to *digg.com* but holds a universal feature regarding phenomena of collective attention.

This pattern can be explained as follow: in the beginning a story attracts the attention of a few people, which increases its popularity. As it becomes popular, the item attracts more and more attention. This triggers a self-reinforcing loop leading to a non-linear increase of attention over time. More precisely, the origin of this burst of attention is twofold: on the one hand, it can result from external influences, such as television broadcasting, and affects a large number of people at the same time. This leads to a sharp response of the social system to the occurrence of the event, as it can be observed for the Asian tsunami. On the other hand, an *endogenous* growth is generated by interactions between people, such as word-of-mouth, and characterized by a progressive rather than sudden increase of attention [5]. In most cases, however, both exogenous and endogenous factors affect the dynamics in variable proportions [9]. Regarding the examples displayed in Fig.1a, the exogenous influence is dominating, either because the mass media have largely relayed the information for the Asian tsunami or because the story appeared in the popular section of digg.com.

As time goes by, however, various limiting factors may play a role, such as the age of the news or the occurrence of other events, making people to turn their attention away from previous news toward novel ones. This finally results in a slow decay of popularity.

[1] *'How digg works'*: http://digg.com/how

[2] '*Google trends*': http://www.google.com/trends

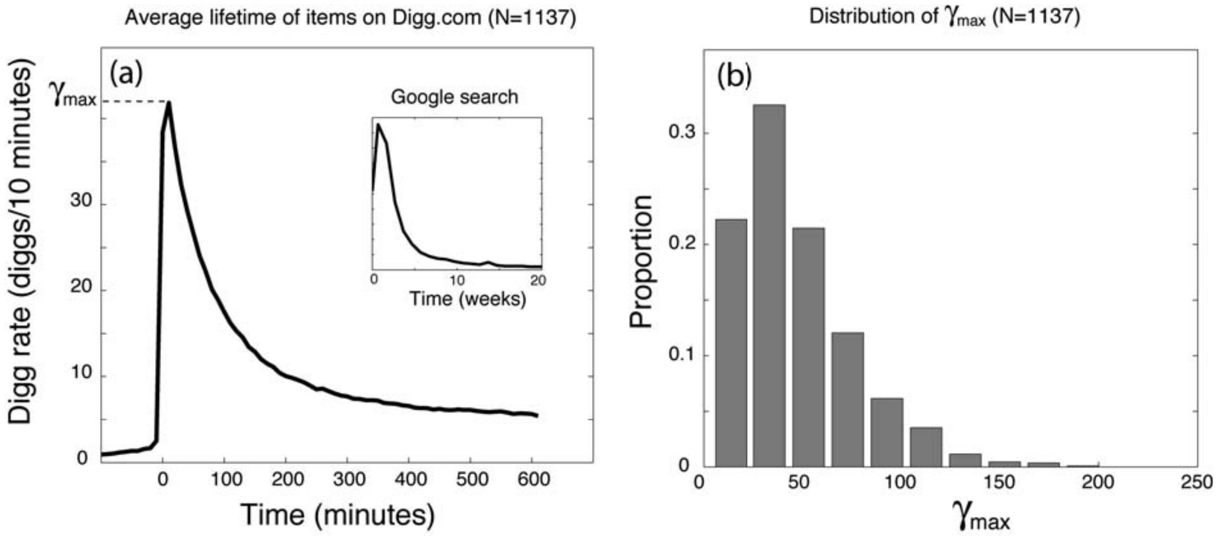


**Figure 1.** (a) Average amount of attention given to stories on digg.com over time. Time t=0 is the time at which the story appeared in the popular section. The inset indicates the amount of search queries for the keyword '*tsunami'* on *google* over time, after the Asian tsunami happened in December 2004. The value $\gamma_{\max}$ indicates the maximum popularity level that the news item has reached. (b) The distribution of the values of $\gamma_{\max}$ for the same dataset.

The study of the dynamics at work on the website *digg.com* actually provides some additional information to characterize features of collective attention. On the basis of the dataset we captured from the website, we focused on the differentiation between the items, that is on how much the maximum level of popularity differed between the stories. Figure 1b shows the distribution of the maximum digg rate $\gamma_{\max}$ each story reached. It turns out that most stories are barely paid attention to, while a few of them attracted the attention of a large number of users. A Kolmogorov-Smirnov normality test of $\log(\gamma_{\max})$ with mean 3.66 and standard deviation 0.64 yields a P-value of 0.0413, suggesting that $\gamma_{\max}$ follows a log-normal distribution. In the following, we consider this statistical feature as a signature of the phenomenon and use it to test and validate our model, which we describe in the following section.

## 3 Model description

In this section, we propose a model of collective attention based on individual behavior. As a starting point, it is natural to ask how attention works at the level of a single individual. In the field of human psychology and neurobiology, it is well-known that the brain has a limited rate of information processing. In particular, one person can either *select* a particular item of the environment for detailed analysis and completely ignore other items, or it can *share* the attention among several targets while missing the details. This corresponds to the so-called *selective* and *divided* attention regime, respectively [10].

Accordingly, we suggest that each agent $j$ of our model has a limited attention capacity $C_j$ that can be shared over many items. The agent can, for example, pay 100% of its attention to a single event, or 70% to a first one and 10% to three others, and so forth. The value of $C_j \geq 0$ is initially set to 1 for all agents.

The environment in which agents interact is defined as a collection of N items, each item $i$ being characterized by its age $A_i$ and popularity $P_i$. The age of each item naturally increases in time, while the popularity is the amount of attention that the item receives from the agents. Once an item reaches a certain maximum age $A_{\max}$, it is removed. The value of $A_{\max}$ must be chosen sufficiently large so that it does not affect the dynamics.

Given this framework, we rely on earlier empirical observations to define three behavioral rules that describe the agents' behavior:

1. Agents are attracted by popular items.
2. Agents spend a certain fraction of their attention capacity on these items.
3. Agents continuously recover a small amount of attention they previously paid to each items.

In the following, we discuss each of these three rules in detail.

**Rule 1.** The first behavioral rule defines how the actors of the model are attracted by the news items. It describes the fact that, in real life, the probability of an uninformed individual to hear about a novelty increases with its popularity.
More specifically, early experiments by Stanley Milgram show that the probability to pay attention to an event strongly increases with the amount of people attending it, but reaches a certain high saturation level [2].
To formalize this idea, we assign to each item $i$ a weight $W_i(t)$ that describes its attractiveness**.** The greater the weight of an item, the higher its probability to be chosen. $W_i(t)$ is defined according to Milgram's observations (see Fig. 2a):

$$W_i(t) = 1 - e^{-k_1 P_i(t)} + \varepsilon$$

where the parameter $k_1$ describes how strong is the effect of increasing popularity. The above equation, however, assigns a weight $W_i = 0$ for items with popularity $P_i = 0$, implying that new and previously unknown items would be ignored by the agents. It is therefore necessary to introduce a fluctuation term $\varepsilon$, that allows the agents to discover unknown items. We defined the exploration rate $\varepsilon$ as a uniformly distributed random value selected in the range 0 to $\varepsilon_0$.
At the end of this first step, each agent $j$ chooses an item, with a probability

$$p_i(t) = \frac{W_i(t)}{\sum_{n=1}^{N} W_n(t)} .$$

**Rule 2.** Once the agent has selected an item, it spends a fraction $\beta_i$ of its attention capacity to it. A large value of $\beta_i$ means that the agent is particularly interested in this item, while a small value implies that the item will be checked out in a superficial way. A recent study based on the website *digg.com* shows that, as an item becomes older, it attracts less attention [6]. This so-called *decay of novelty* follows a stretched exponential law with increasing age. In our model, we describe the fraction $\beta_i$ as a function of the age $A_i$ of a news item by:

$$\beta_i(t) = e^{-k_2 A_i(t)}$$

That is, we replaced the stretched exponential by an exponential law for simplicity (see Fig. 2b). The amount $\beta_i$ is then subtracted from the attention capacity $C_j$ of the agent and added to the popularity $P_i$ of the item. If the corresponding $\beta_i$ would exceed the remaining capacity $C_j$ of the agent, the value $C_j$ is taken instead (so that the attention capacity of the agents can never fall below zero).

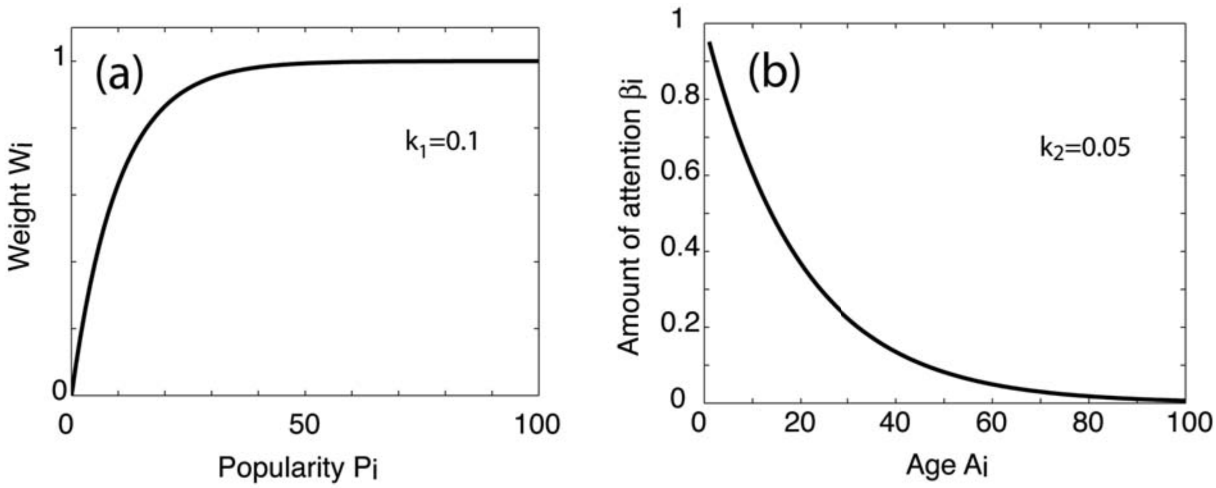


**Figure 2.** Specification of our individual-based model. (a) The weight of an item as a function of its popularity. This dependency follows from Milgram's experiments in Ref. [2]. (b) The amount of attention an agent gives to an item as a function of the age of the item. The exponential decay is a slightly simplified specification of the empirically observed dependency described in Ref. [6].

**Rule 3.** Finally, we consider that agents continuously recover a small amount $\alpha$ of attention they previously gave to each item. For all items $i$ currently attended by agent $j$, the amount $\alpha$ is deduced from popularity $P_i$ and goes back to its attention capacity $C_j$. As for the previous rule, if $\alpha$ exceeds the amount of attention agent $j$ is giving to item $i$, this latter value is taken instead, so that the total amount of attention in the system remains constant.

In principle, the value of $\alpha$ reflects how interesting a news item is. High values of $\alpha$ would make the agent quickly turn away from the item, while for low values, the agent attend the item for a longer period of time. As a first approximation, we set $\alpha$ constant for all items, regardless of their quality. We note, however, that it would be interesting to specify $\alpha_i$ as a function of the relevance of item $i$. This would allow one to compare the evolution of interesting news in comparison with non-interesting ones.

## 4 Numerical investigation of the model

Given the above model, we will now investigate how well simulations fit the observed patterns previously described, and what we can learn regarding the phenomena under study.
For this, we have conducted a set of simulations based on the three rules describes above. The population was set to n=100 agents, each having an initial attention capacity $C_j = 1$. At each time step, a news item is added with probability $p_{in}$ describing the inflow of information. At time t=0, ten items were arbitrarily introduced, and a random amount of attention was transferred from the agents' attention capacities to these items. To avoid biases to the initial conditions, these 10 items were left out from the evaluation of the simulation data.

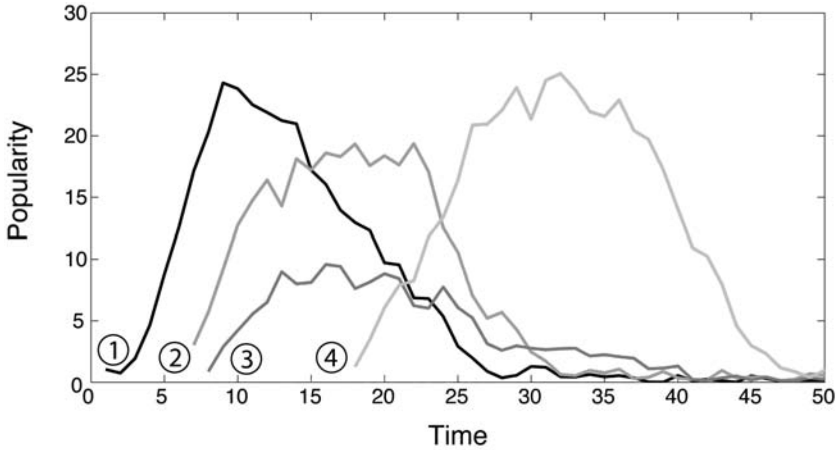


**Figure 3.** Sequence of four news items appearing successively during a typical simulation run.

Figure 3 displays a short sequence of a typical simulation run, showing the evolution of popularity for four successively appearing news items. As expected, novel items are eventually discovered by agents at random a short time after their occurrence and attract a variable amount of attention from the agents. Because of their increasing age, the popularity level of older items decreases, and they vanish some time later. This dynamics yields an average curve of popularity that is similar to the observed ones (see Fig. 4a).
For comparison with the dynamics observed on *digg.com*, we have defined the variable $P_{\max}$ as the maximum attention that an item has attracted during its lifetime (i.e. the peak value of the attention burst). The distribution of $P_{\max}$ follows, in fact, a log-normal distribution (Pvalue=0.23). This result is in good agreement with observations described in section 2 and reflects the fact that most items are barely considered by the agents, while a few of them reach a high level of popularity (Fig. 4b).

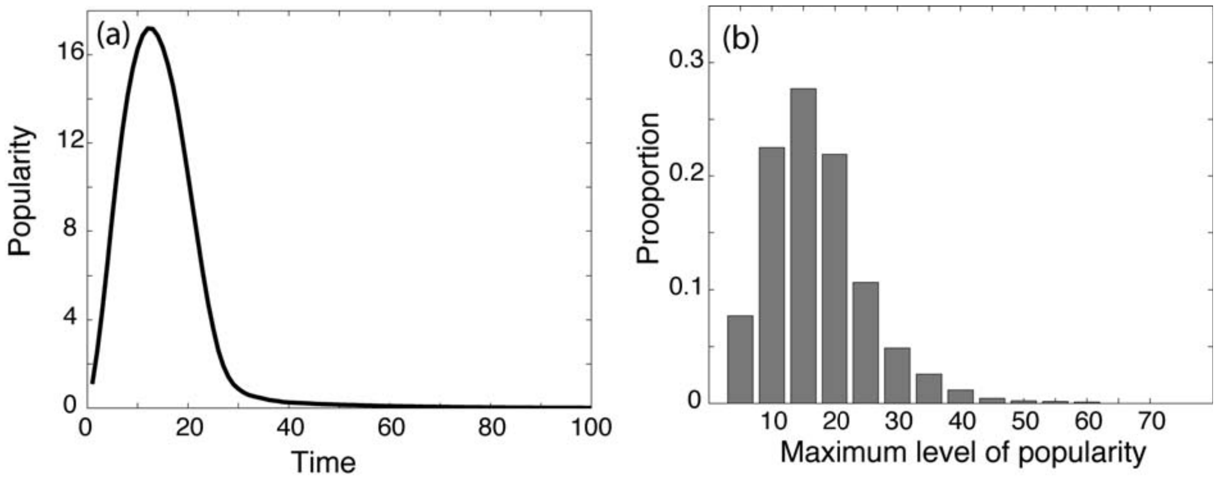


**Figure 4.** Simulation results. (a) Average popularity dynamics of the item. (b) Distribution of the maximum level of popularity. The parameters used in our simulation are as follows: n=100 agents; $p_{in}$=0.25; $\varepsilon_0$=0.2; $\alpha$=0.15; $A_{\max}$=100; $k_1$=0.1; $k_2$=0.05

The simulation sequence presented in figure 3 also highlights an interesting effect of competition between items. It seems that a news item attracts a lower amount of attention if it appears at a time when many others are already present (for example, compare the popularity of the third and fourth item). Because the total amount of attention remains constant in time, any item that becomes very popular tends to deprive other news items from a significant part of the population's attention.

To investigate this effect further, we have measured the cumulated amount of attention $P_{cumul}$ each item attracted during its lifetime as a function of the news inflow $p_{in}$. The model predicts a decay of the average popularity with an increasing flow of information (see figure 5a). Further on, we have investigated whether this competition effect is also visible on digg.com. For this, we have determined the correlation between the average final number of diggs of stories and the number of concurrent stories posted within the previous or following 30 minutes. For all 37316 popular stories posted between July 2007 and June 2008, we found that the average number of diggs a story reached decreases with the number of concurrent stories (Fig. 5b). An exponential decay fits the observed dataset well (R-square=0.84), in accordance with the model prediction. Interestingly, this dynamics has been observed on digg.com regardless of the news content, topic or relevance, suggesting that the quality of the item would often have a minor influence on the collective dynamics as compared to the timing of its release.

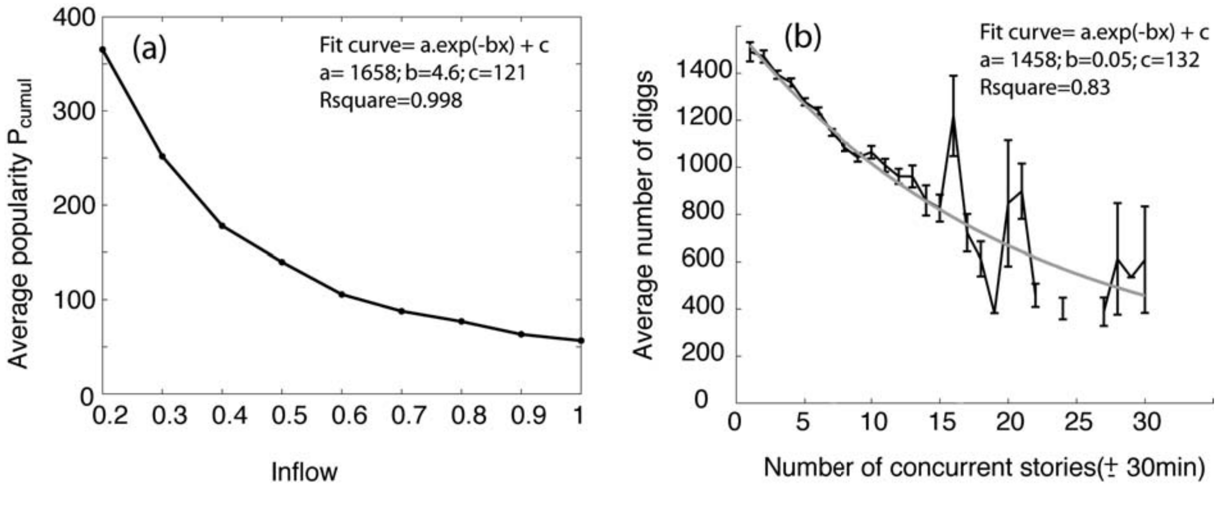


**Figure 5.** Effect of competition between concurrent news. (a) The simulation model predicts the average amount of attention for each news item decreases with increasing arrival of news. (b) The average final number of diggs for 37316 stories on digg.com as a function of the number of concurrent stories posted within the same period of time. Error bars indicate the standard error of the mean value.

## 5 Discussion

In this article, we have presented a simple individual-based model of collective attention. The assumed model ingredients were mainly based on empirical observations. Our simulation results display several encouraging similarities with the dynamics observed on the website *digg.com*. In a situation where 100 agents were facing a flow of new information, the model predicted a burst of attention around novel items and the log-normal distribution of popularity of different news items.
Our model predicted that the expected popularity of a news item strongly depends on the number of items appearing approximately at the same time. This could be confirmed by data from digg.com. We expect this as an effect of competition for a limited attention capacity. Our observations may be compared with situations where several commercial products (e.g. movies or books) are released at almost the same time [11].
An interesting future research perspective concerns the effect of the initial weight of news items. So far, we have assumed that incoming items would have a random chance to be discovered by the agents. Assigning a higher initial weight to an item comes down to increasing its attractiveness regardless of its current popularity, similarly to an exogenous origin of the event. This parameter therefore allows one to reflect the external influence of a news item and may be used to investigate how mass media influence the behavior of people in an information-rich environment.

It is interesting to note that the suggested model describes a characteristic dynamics of collective behavior, which can be observed in a variety of social and biological systems [12][13]. The combination of positive feedbacks (amplification loops) and negative feedbacks (relaxation effects) are typical ingredients that lead a group to self-organize and form collective patterns.
For example, the suggested model relates to various other studies in which people own a limited amount of 'energy' that may be dynamically distributed over several items. In the study of social networks, for example, it has been demonstrated that such a mechanism is enough to understand the formation of friend networks, where individuals have limited friendship capacities that they share dynamically among their friends [14].
In biological systems, the phenomenon of collective attention is also similar in many points to the process that leads ant colonies to select a foraging option among many potential alternatives [15]. Interestingly, one major difference is that ants will make a collective choice for a unique foraging option and concentrate all the activity of the colony on it, while people rather tend to explore all the alternatives and share their attention over the items in various proportions (potentially spending a longer time on a few specific elements). Our study, therefore, provides some new elements to compare and better understand various aspects of collective behaviors in social and biological systems.

## 6 References


[1] Simon, H.A., Designing Organizations for an Information-Rich World, in: Baltimore, MD, in Martin Greenberger, Computers, Communication, and the Public Interest, The Johns Hopkins Press (1971)

[2] Milgram, S., Bickman, L., Berkowitz, L., Note on the drawing power of crowds of different size. Journal of Personality and Social Psychology, 13 (1969), 79-82.

[3] Börner, K., Maru, J. T., Goldstone, R. L. The simultaneous evolution of author and paper networks. Proceedings of the National Academy of Science, 101, (2004), 5266-5273.

[4] Schweitzer, F., Mach, R., The epidemics of donations: Logistic growth and power-laws. PLoS ONE, 3(1), 2008

[5] Crane, R., Sornette, D., Robust dynamic classes revealed by measuring the response function of a social system. Proceedings of the National Academy of Science, 105(41), (2008), 15649-15653.

[6] Wu, F., Huberman, B.A., Novelty and collective attention, Proceedings of the National Academy of Science, 104(45), (2007), 17599-17601

[7] Wilkinson D.M., Huberman B.A., Assessing the Value of Coooperation in Wikipedia, First Monday, 12(4) 2007.

[8] Golder S.A., Wilkinson D. & Huberman B.A., Rhythms of social interaction: Messaging within a massive online network. 3rd International Conference on Communities and Technologies (CT2007). East Lansing, MI. June 28-30, 2007.

[9] Lai G., Wong O., The tie effect on information dissemination: The spread of a commercial rumor in Hong Kong, Social Networks 24 (2002), 49-75.

[10] Corbetta M., Miezin F.M., Dobmeyer S.M., Shulman G.L. & Petersen S.E. Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. The Journal of Neuroscience, 11(8):2383-2402 (1991)

[11] Falkinger, J., Attention economies, Journal of Economic Theory, 133, (2007), 266-294
[12] Schelling, T.C., Micromotives and Macrobehavior. (1978) New York, Norton.
[13] Camazine, S., Deneubourg, J.L., Franks, N.R., Sneyd, J. Theraulaz, G. Bonabeau, E. Self-organization in Biological Systems (2001) Princeton University Press, Princeton, NJ.
[14] Whitmeyer, J.M., Yeingst C.M., A dynamic model of friendly association networks, Social Science Research 35 (2006) 642–667
[15] Deneubourg J.L., Goss S., Collective patterns and decision-making, Ethology, Ecology, and Evolution 1 (1989) 295-311