

Recommendation systems in the scope of opinion formation: a model

Marcel Blattner
Laboratory for Web Science
University of Applied Sciences FFHS
Regensdorf, Switzerland
marcel.blattner@ffhs.ch

Matus Medo
Physics Department
University of Fribourg
Fribourg, Switzerland
matus.medo@unifr.ch

ABSTRACT

Aggregated data in real world recommender applications often feature fat-tailed distributions of the number of times individual items have been rated or favored. We propose a model to simulate such data. The model is mainly based on social interactions and opinion formation taking place on a complex network with a given topology. A threshold mechanism is used to govern the decision making process that determines whether a user is or is not interested in an item. We demonstrate the validity of the model by fitting attendance distributions from different real data sets. The model is mathematically analyzed by investigating its master equation. Our approach provides an attempt to understand recommender system's data as a social process. The model can serve as a starting point to generate artificial data sets useful for testing and evaluating recommender systems.

Categories and Subject Descriptors

H.1.1.m [Information Systems]: Miscellaneous

General Terms

Experimentation, Theory

Keywords

recommender systems, opinion formation, complex networks

1. INTRODUCTION

This is the information age. We are witnessing information production and consumption in a speed never seen before. The WEB2.0 paradigm enables consumers and producers to exchange data in a collaborative way benefiting both parties. However, one of the key challenges in our digitally-driven society is information overload [7]. We have the 'pain of choice'. Recommendation systems represent a possible solution to this problem. They have emerged as a research area

Paper presented at the 2012 Decisions@RecSys workshop in conjunction with the 6th ACM conference on Recommender Systems. Copyright 2012 for the individual papers by the papers' authors. Copying permitted for private and academic purposes. This volume is published and copyrighted by its editors.

on its own in the 90s [42, 20, 28, 21, 11]. The interest in recommendation systems increased steadily in recent years, and attracted researchers from different fields [43]. The success of highly rated Internet sites as Amazon, Netflix, YouTube, Yahoo, Last.fm and others is to a large extent based on their recommender engines. Corresponding applications recommend everything from CD/DVD's, movies, jokes, books, web sites to more complex items such as financial services.

The most popular techniques related to recommendation systems are collaborative filtering [8, 26, 11, 24, 28, 21, 41, 45] and content-based filtering [14, 40, 35, 5, 30]. In addition, researchers developed alternative methods inspired by fields as diverse as machine learning, graph theory, and physics [16, 17, 37, 52, 51, 10, 48, 50]. Furthermore, recommendation systems have been investigated in connection with trust [2, 39, 47, 32, 33] and personalized web search [9, 12, 46], which constitutes the new research frontier in search engines.

However, there are still many open challenges in the research field of recommendation systems [1, 22, 25, 18, 24, 43, 15]. One key question is connected to the understanding of the user rating mechanism. We build on a well documented influence of social interactions with peers on the decision to vote, favor, or even purchase an item [44, 27]. We propose a model inspired by opinion formation taking place on a complex network with a predefined topology. Our model is able to generate data observed in real world recommender systems. Despite its simplicity, the model is flexible enough to generate a wide range of different patterns. We mathematically analyze the model using a mean field approach to the full Master Equation. Our approach provides an understanding of the data in recommender systems as a product of social processes. The model can serve as a data generator which is valuable for testing and evaluation purposes for recommender systems.

The rest of the paper is organized as follows. The model is outlined in Sec. (2). Methods, data set descriptions, and validation procedures are in Sec. (3). Results are presented in Sec. (4). Discussion and an outlook for future research directions are in Sec. (5).

2. MODEL

2.1 Motivation

Our daily decisions are heavily influenced by various information channels: advertisement, broadcastings, social interactions, and many others. Social ties (word-of-mouth) play a pivotal role in consumers buying decisions [44, 27]. It was

demonstrated by many researchers that personal communication and informal information exchange not only influence purchase decisions and opinions, but shape our expectations of a product or service [49, 4, 3]. On the other hand, it was shown [23], that social benefits are a major motivation to participate on opinion platforms. If somebody is influenced by recommendations on an opinion platform like MovieLens or Amazon, social interactions and word-of-mouth in general are additional forces governing the decision making process to purchase or even to rate an object in a particular way [31].

Our model is formulated within an opinion formation framework where social ties play a major role. We shall discuss the following main ingredients of our model:

- Influence-Network (IN)
- Intrinsic-Item-Anticipation (IIA)
- Influence-Dynamics (ID)

Influence Network.

We call the network where context-relevant information exchange takes place an Influence-Network (IN). Nodes of the IN are people and connections between nodes indicate the influence among them. Note that we put no constraints on the nature of how these connections are realized. They may be purely virtual (over the Internet) or based on physical meetings. We emphasize that INs are domain dependent, i.e., for a given community of users, the Influence Network concerning books may differ greatly (in topology, number of ties, tie strength, etc.) from that concerning another subject such as food or movies. Indeed, one person's opinion leaders (relevant peers) concerning books may be very different from those for food or other subjects. In this scope, we see the INs as domain-restricted views on social networks. It is thus reasonable to assume that Influence Networks are similar to social interaction networks which often exhibit a scale-free topology [6]. However, our model is not restricted to a particular network structure.

Intrinsic-Item-Anticipation.

Suppose a new product is launched on the market. Advertisement, marketing campaigns, and other efforts to attract customers predate the launching process and continue after the product started to spread on the market. These efforts influence product-dependent customer anticipation. It is clear that the resulting anticipation is a complex combination of many different components including intrinsic product quality and possibly also suggestions from recommendation systems.

In our model we call the above-described anticipation Intrinsic-Item-Anticipation (IIA) and measure it by a single number. It is based on many external sources, except for the influence generated by social interactions. It is the opinion on something taken by individuals, before they start to discuss the subject with their peers. Furthermore, we assume that an individual will invest resources (time/money) into an object only, if the Intrinsic-Item-Anticipation is above a particular threshold, which we call Critical-Anticipation-Threshold.

Influence-Dynamics.

The Influence-Dynamics describes how individuals' Intrinsic-Item-Anticipations are altered by information exchange via

the connections of the corresponding Influence-Network. From our model's point of view this means the following: an individual's IIA for a particular item i may be shifted due to social interactions with directly connected peers (these interactions thus take place on the corresponding IN), who already experienced the product or service in question. This process can shift the Intrinsic-Item-Anticipation of an individual who did not yet experience product/object i closer to or beyond the critical-anticipation-threshold.

We now summarize the basic ingredients of our model. An individual user's opinions on objects are assembled in two consecutive stages: i) opinion making based on different external sources, including suggestions by recommendation systems and ii) opinion making based on social interactions in the Influence-Network. The second process may shift the opinions generated by the first process.

2.2 Mathematical formulation of the model

In this section we firstly describe how individuals' Intrinsic-Item-Anticipations may change due to social interactions taking place on a particular Influence-Network. Secondly, we introduce dynamical processes governing the opinion propagation.

IIA shift.

We model a possible shift in the IIA as:

$$\hat{f}_{ij} = f_{ij} + \left[\frac{\Theta_j}{k_j} \right]^{(1-\gamma)}. \quad (1)$$

where \hat{f}_{ij} is the shifted Intrinsic-Item-Anticipation of individual j for object i , f_{ij} is the unbiased IIA, Θ_j is the number of j 's neighbors, who already experienced and liked item i , k_j denotes the total number of j 's neighbors in the corresponding IN, and $\gamma \in (0, 1)$ quantifies trust of individuals to their peers. An individual j will consume, purchase, or positively rate an item i only if

$$\hat{f}_{ij} \geq \Delta. \quad (2)$$

We identify Δ as the Critical-Anticipation-Threshold. Values of f_{ij} are drawn from a probability distribution f_i . Since the IIA for each individual is an aggregate of many different and largely independent contributions, we assume that f_i is normally distributed, $f_i \in \mathcal{N}(\mu_i, \sigma)$. (Unless stated otherwise.) To mimic different item anticipations for different objects i , we draw the mean μ_i from a uniform distribution $U(-\epsilon, \epsilon)$. We maintain μ_i , ϵ , and σ , so that f_i is roughly bounded by $(-1, 1)$, i.e., $-1 \leq \mu - 3\sigma < \mu + 3\sigma \leq 1$. Note that \hat{f}_{ij} can exceed these boundaries after a shift of the corresponding IIA occurs. The second term on the right hand side of Eq. (1) is the influence of j 's neighborhood weighted by trust γ . To better understand the interplay between γ and the density of attending users in the neighborhood of user i , $\rho := \Theta_j/k_j$, we refer to Fig. 1. Trust $\gamma \approx 1$ causes a big shift on the IIA's even for $\rho \approx 0$. On the other hand, $\gamma \approx 0$ needs high ρ to yield a significant IAA shift. These properties are understood as follows: people trusting strongly in their peers need only few positive opinions to be convinced, whereas people trusting less in their social environment need considerable more signals to be influenced.

Influence-Dynamics.

The Influence-Dynamics proceeds as follows. Firstly, we

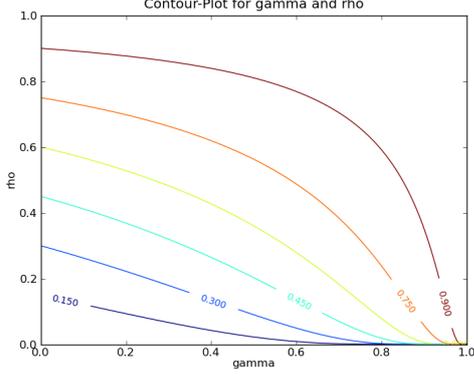


Figure 1: Contour plot for γ and $\rho = \Theta_j/k_j$. Numbers inside the plot quantify the shift in the IAA as a function of γ and ρ .

draw an Influence-Network $IN(\mathcal{P})$ with a fixed network topology (power-law, Erdős-Rényi, or another). \mathcal{P} refers to a set of appropriate parameters for the Influence-Network in question (like network type, number of nodes, etc.). The network's topology is not affected by the dynamical processes (opinion propagation) taking place on it. We justify this static scenario by assuming that the time scale of the topology change is much longer than the time scale¹ of opinion spreading in the network. Each node in the Influence-Network corresponds to an individual. For each individual j we draw an unbiased Intrinsic-Item-Anticipation f_{ij} from the predefined probability distribution f_i . At each time step, every individual is in one of the following states: $\{S, A, D\}$. S refers to a susceptible state and corresponds to the initial state for all nodes at $t = 0$. A refers to an attender state and corresponds to an individual with the property $\hat{f}_{ij} \geq \Delta$. D refers to a denier state with the property $\hat{f}_{ij} < \Delta$ after an information exchange with his/her peers in the Influence-Network happened. An individual in state D or A can not change his/her state anymore. It is clear that an individual in state A cannot back transform to the susceptible state S , since he/she did consume or favor item i and we do not account for multiple attendances in our model. An individual in state D was influenced but not convinced by his opinion leaders (directed connected peers). We make the following assumption here: if individual j 's opinion leaders are not able to convince individual j , meaning that individual's j Intrinsic Item Anticipation \hat{f}_{ij} stays below the critical threshold Δ after the influence process, then we assume that j 's opinion not to attend object i remains unchanged in the future. Therefore we have the following possible transitions for each node in the influence network: $j_S \rightarrow j_A$ or $j_S \rightarrow j_D$. Node states are updated asynchronously which is more realistic than synchronous updating, especially in social interaction models [13]. The Influence-Dynamics is summarized in Algorithm 1.

Master Equation.

We are now in the position to formulate the Master Equa-

¹The term time scale denotes a dimensionless quantity and specifies the deviations of time. A shorter time scale means a faster spreading of opinions in the network.

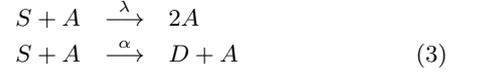
Algorithm 1 RecSysMod algorithm. \mathcal{P} contains the configuration parameter for the network. Δ is the Anticipation Threshold and γ denotes the trust. $O \in \mathbb{N}$ is the number of objects to simulate. $G(N, E)$ is the network. N is the set of nodes and E is the set of edges.

```

1: procedure RECSYSMOD_I( $\mathcal{P}, \Delta, \gamma, O$ )
2:    $G(N, E) \leftarrow \text{GenNetwork}(\mathcal{P})$ 
3:   for all Objects in  $O$  do
4:     generate distribution  $f_i$  from  $\mathcal{N}(\mu_i, \sigma)$ 
5:     for each node  $j \in N$  in  $G$  do
6:       draw  $f_{ij}$  from  $f_i$ 
7:       if  $f_{ij} < \Delta$  then
8:          $j_{state} \leftarrow S$ 
9:       else
10:         $j_{state} \leftarrow A$ 
11:      end if
12:    end for
13:  repeat
14:    for all  $j$  with  $j_{state} = S$  AND  $\Theta_j > 0$  do
15:       $\hat{f}_{ij} \leftarrow f_{ij} + \left[ \frac{\Theta_j}{k_j} \right]^{(1-\gamma)}$ 
16:      if  $\hat{f}_{ij} < \Delta$  then
17:         $j_{state} \leftarrow D$ 
18:      else
19:         $j_{state} \leftarrow A$ 
20:      end if
21:    end for
22:  until  $|\{j | j_{state} = S \text{ AND } \Theta_j > 0\}| = 0$ 
23: end for
24: end procedure

```

tion for the dynamics. As already said before, two things can happen when a non-attender is connected to an attender: a) he/she becomes an attender too, or b) he/she becomes a denier who will not attend/favor the item. For these two interaction types we formally write:



Here λ denotes the probability that a susceptible node connected to an attender becomes an attender too, and α is the probability that a susceptible node attached to an attender becomes a denier. To take into account the underlying network topology of the Influence Network it is common to introduce compartments k [19]. Let N_k^A be the number of nodes in state A with k connections, N_k^S the number of nodes in state S with k connections, and N_k^D the number of nodes in state D with k connections, respectively. Furthermore we define the corresponding densities: $a_k(t) = N_k^A/N_k$, $s_k(t) = N_k^S/N_k$ and $d_k(t) = N_k^D/N_k$. N_k is the total number of nodes with k connections in the network. Since every node from N_k must be in one of the three states, $\forall t : a_k(t) + s_k(t) + d_k(t) = 1$. A weighted sum over all k compartments gives the total fraction of attenders at time t , $a(t) = \sum_k P(k)a_k(t)$ where $P(k)$ is the degree distribution of the network (it also holds that $a(t) = N^A(t)/N$). The time dependence of our state variables $a_k(t), d_k(t), s_k(t)$ is

$$\left. \begin{aligned}
 \dot{a}_k(t) &= \lambda k s_k(t) \Omega \\
 \dot{d}_k(t) &= \alpha k s_k(t) \Omega \\
 \dot{s}_k(t) &= -(\alpha + \lambda) k s_k(t) \Omega
 \end{aligned} \right\} \tag{4}$$

where Ω is the density of attenders in the neighborhood of susceptible node with k connections averaged over k

$$\Omega = \sum_k P(k)(k-1)a_k / \langle k \rangle \quad (5)$$

where $\langle k \rangle$ denotes the mean degree of the network. As outlined above, λ is the probability that a node in state S transforms to state A if it is connected to a node in state A . This happens when $\hat{f}_{ij} > \Delta$. Therefore, we have $\Delta_- < f_{ij} < \Delta$ where $\Delta_- = \Delta - (1/k)^{1-\gamma}$. From this we have $\lambda = \int_{\Delta_-}^{\Delta} f(x)dx$, where $f(x)$ is the expectation distribution. Similarly we write for $\alpha = \int_l^{\Delta_-} f(x)dx$, where l denotes the lower bound of the expectation distribution $f(x)$. A crude mean field approximation can be obtained by multiplying the right hand sides of Eq. (4) with $P(k)$ and summing over k , which yields a set of differential equations

$$\left. \begin{aligned} \dot{a}(t) &= \lambda \langle k \rangle s(t)a(t), \\ \dot{d}(t) &= \alpha \langle k \rangle s(t)a(t), \\ \dot{s}(t) &= -(\alpha + \lambda) \langle k \rangle s(t)a(t). \end{aligned} \right\} \quad (6)$$

which is later used to obtain analytical results for the attendance fraction $a(t)$.

3. METHODS

We describe here our simulation procedures, datasets, experiments, and analytical methods.

Simulations.

Our simulations employ Alg. (1). As outlined in the model section, we do not change the network topology during the dynamical processes. We experiment with two different network types, Erdős-Rényi (ER), and power law (PL) which are both generated by a so-called configuration model [34]. ER and PL represent two fundamentally different classes of networks. The former is characterized by a typical degree scale (mean degree of the network), whereas the latter exhibits a fat-tailed degree distribution which is scale free. The networks are random and have no degree correlations and no particular community structure. To obtain representative results we stick to the following approach: we fix the network type, number of nodes, number of objects, and network type relevant parameters to draw an ER or PL network. We call this a configuration \mathcal{P} . In addition, we fix the variance σ of the anticipation distributions f_i . We perform each simulation on 50 different networks belonging to the same configuration \mathcal{P} and on each network we simulate the dynamics 50 times. Then we average the obtained attendance distributions over all 2500 simulations.

Datasets.

To show the validity of our model we use real world recommender datasets. **MovieLens** (movielens.umn.edu), a web service from GroupLens (grouplens.org) where ratings are recorded on a five stars scale. The data set contains 1682 movies and 943 users. Only 6,5% of possible votes are expressed. **Netflix** data set (netflix.com). We use the Netflix grand prize data set which contains 480189 users and 17770 movies and also uses a five stars scale. **Lastfm** data set (Lastfm.com). This data set contains social networking, tagging, and music artist listening information from users of the Last.fm online music system. There are 1892 users,

17632 artists, and 92834 user-listened artists relations in total. In addition, the data set contains 12717 bi-directional user friendship relations. These data sets are chosen because they exhibit very different attendance distributions and thus provide an excellent playground to validate our model in different settings.

Experiments.

Data topologies. We firstly investigate the simulated attendance distributions as a function of trust γ , the anticipation threshold Δ , and the network topology. For this purpose we simulate the dynamics on a toy network with 500 nodes and record the final attendance number of 300 objects. The simulation is conducted for ER and PL networks and performed as outlined in the simulations paragraph above. In Fig.(2) and Fig.(3) we investigate the skewness [53] of the attendance distributions and the maximal attendance obtained for the corresponding parameter settings. The skewness of a distribution is a measure for the asymmetry around its mean value. A positive skewness value means that there is more weight to the left from the mean, whereas a negative value indicates more weight in the right from the mean.

Fitting real data. We explore the model's ability to fit real world recommendation attendance distributions found in the described data sets. For this purpose we fix for the Netflix data set a network with 480189 nodes and perform a simulation for 17770 objects. In the MovieLens case we do the same for 943 nodes and 1682 objects and for the Lastfm data set we simulate on a network with 1892 nodes and 17632 objects. In the case of Lastfm we have the social network data as well. We validate our model on that data set by two experiments: a) we use the provided user friendship network as simulation input and fit the attendance distribution and b) we fit the attendance distribution like in the MovieLens and Netflix case with an artificially generated network.

Mathematical analysis. We investigate the Master Equations Eq. (4) and Eq. (6). We provide a full analytical solution for Eq. (6) and an analytical approximation for Eq. (4) in the early spreading stage.

4. RESULTS

Data topologies. The landscape of attendance distributions of our model is demonstrated in Fig. (2) and Fig. (3). To obtain these results, simulations were performed as described in Sec. (3). The item anticipation f_i was drawn from a normal distribution with mean values $\mu_i \in U(-0.1, 0.1)$ and variance $\sigma = 0.25$ fixed for all items. Both networks have 500 nodes. In the Erdős-Rényi case, we used a wiring probability $p = 0.03$ between nodes. The Power Law network was drawn with an exponent $\delta = 2.25$. The simulated attendance distributions in Fig.(2) and Fig.(3) show a wide range of different patterns for both ER and PL Influence-Networks. In particular, both network types can serve as a basis for attendance distributions with both positive and negative skewness. Therefore, the observed fat-tailed distributions are not a result of the heterogeneity of a scale free network but they are emergent properties of the dynamics produced by our model. The parameter region for highly positively-skewed distributions is the same for both network types. The parameters γ and Δ can be tuned so that all items are attended by everybody or all items are attended by nobody. While not relevant for simulating realistic attendance distributions, these extreme cases help to understand

the model's flexibility.

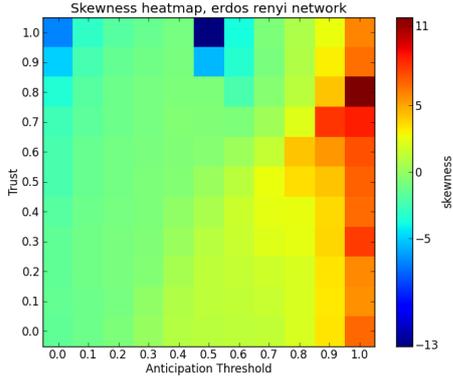


Figure 2: Skewness of the attendance distributions as a function of trust γ and the critical anticipation threshold Δ for Erdős-Rényi networks with 500 nodes and 300 simulated items.

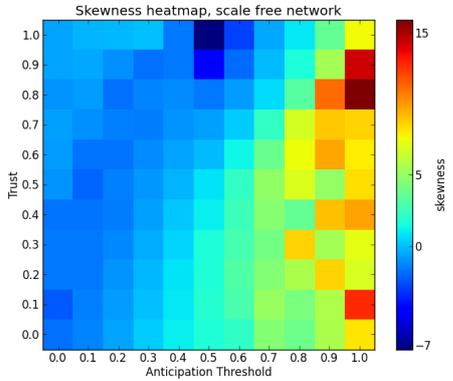


Figure 3: Skewness of the attendance distributions as a function of trust γ and the critical anticipation threshold Δ for power-law networks with 500 nodes and 300 simulated items.

Fitting real data We fit real world recommender data from MovieLens, Netflix and Lastfm with results reported in Fig. (4), Fig. (5), Fig. (6), Fig. (7), and Tab. (1), respectively. The real and simulated distributions are compared using Kullback-Leibler (KL) divergence [29]. We report the mean, median, maximum, and minimum of the simulated and real attendance distributions. Trust γ , anticipation threshold Δ , and anticipation distribution variance σ are reported in figure captions. We also compare the averaged mean degree, maximum degree, minimum degree, and clustering coefficient of the real Lastfm social network and networks obtained to fit the data. Results are reported in Tab. (2) and Fig. (8). Note that thus obtained parameter values can be useful also in real applications where, assuming that our social opinion formation model is valid, one could detect decline of the overall trust value in an online community, for example.

Mathematical analysis. Eq. (6) can be solved analytically. We have $\forall t : a(t) + s(t) + d(t) = 1$ with the initial con-

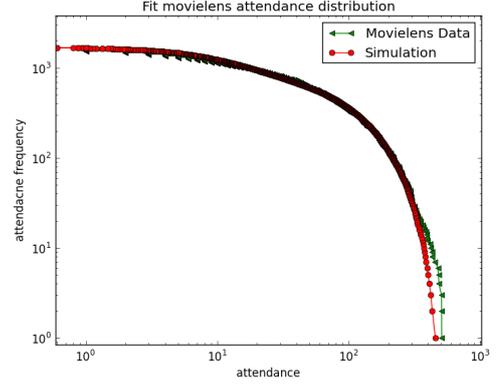


Figure 4: Fit of the MovieLens attendance distribution with trust $\gamma = 0.50$, critical anticipation threshold $\Delta = 0.6$, anticipation distribution variance $\sigma = 0.25$, and power law network with exponent $\delta = 2.25$, 943 nodes, and 1682 simulated objects.

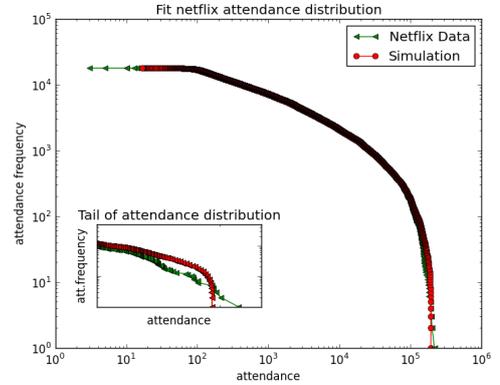


Figure 5: Fit of the Netflix attendance distribution with trust $\gamma = 0.52$, critical anticipation threshold $\Delta = 0.72$, anticipation distribution variance $\sigma = 0.27$, and power law network with exponent $\delta = 2.2$, 480189 nodes, and 17770 simulated objects.

ditions for the first movers $a_0 = \int_{\Delta}^u f(x) dx$, $s(0) = 1 - a(0)$, and $d(0) = 0$. In the following we use the bra-ket notation $\langle x \rangle$ to represent the average of a quantity x . Standard methods can now be used to arrive at²

$$a(t) = \frac{(\tau \langle k \rangle)^{-1} \exp(t/\tau)}{(\alpha + \lambda) [\exp(t/\tau) - 1] + (\tau \langle k \rangle a_0)^{-1}}. \quad (7)$$

Here τ is the time scale of the propagation which is defined as

$$\tau = (a_0 \alpha \langle k \rangle + \lambda \langle k \rangle)^{-1}. \quad (8)$$

This is similar to the time scale $\tau = (\lambda \langle k \rangle)^{-1}$ in the well known SI Model [38, 6]. Eq.(7) can be very useful in predicting the average behavior of users in a recommender system.

Since Eq. (4) is not accessible to a full analytical solution, we investigate it for the early stage of the dynamics. As-

²We give here only the solution for $a(t)$ because we are mainly interested in the attendance dynamics.

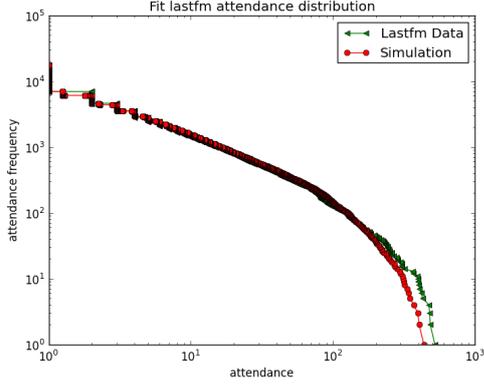


Figure 6: Fit of the Lastfm attendance distribution with trust $\gamma = 0.4$, critical anticipation threshold $\Delta = 0.8$, anticipation distribution variance $\sigma = 0.24$, and real Lastfm user friendship network with 1892 nodes and 17632 simulated objects.

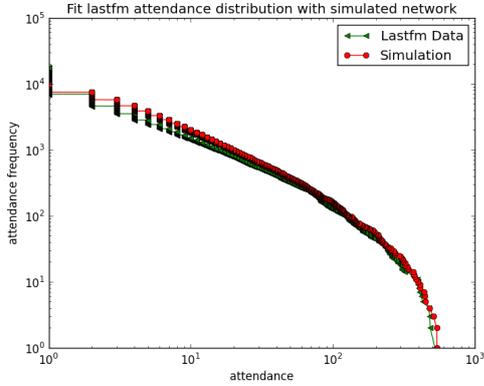


Figure 7: Fit of the Lastfm attendance distribution with trust $\gamma = 0.6$, critical anticipation threshold $\Delta = 0.8$, anticipation distribution variance $\sigma = 0.24$, and power law network with exponent $\delta = 2.25$, 1892 nodes and 17632 simulated objects.

suming $a(0) = a_0 \gg 0$, we can neglect the dynamics of $d(t)$ to obtain

$$\dot{\Omega}(t) = \left(\frac{\langle k^2 \rangle}{\langle k \rangle} - 1 \right) \Omega(t).$$

In addition, Eq. (4) yields

$$\left. \begin{aligned} \dot{a}_k(t) &= \lambda k(1 - a_k(t))\Omega(t) \\ \dot{s}_k(t) &= -(\alpha + \lambda)k(1 - a_k(t))\Omega(t) \end{aligned} \right\} \quad (9)$$

Neglecting terms of order $a_k^2(t)$ and summing the solution of $a_k(t)$ over $P(k)$, we get a result for the early spreading stage

$$a(t) = a(0) \left(1 + \tau \lambda (\exp(t/\tau) - 1) \right), \quad (10)$$

with the timescale $\tau = \langle k^2 \rangle / [\lambda(\langle k^2 \rangle - \langle k \rangle)]$. The obtained time scale τ valid in the early stage of the opinion spreading is clearly dominated by the network heterogeneity. This result is in line with known disease models, e.g., SISIR [38, 6].

D	KL	Med	Mean	Max	Min
ML	0.046	27/26	59/60	583/485	1/1
NF	0.030	561/561	5654/5837	232944/193424	3/16
LFM1	0.05	1/1	5.3/5.2	611/503	1/1
LFM2	0.028	1/1	5.3/5.8	611/547	1/1

Table 1: Simulation results. ML: Movielens, NF: Netflix, LFM1: Lastfm with real network, LFM2: Lastfm with simulated network, KL: Kullback-Leibler divergence, Med: Median, Mean, Max: maximal attendance (data/simulated), Min: minimal attendance (data/simulated).

D	$\langle k \rangle$	k_{min}	k_{max}	δ	C
LFM1	13.4	1	119	2.3	0.186
LFM2	12.0	1	118	2.25	0.06

Table 2: Mean, minimum, maximum degree, clustering coefficient C , and estimated exponent δ of the real (LFM1) and simulated (LFM2) social network for the Lastfm data set.

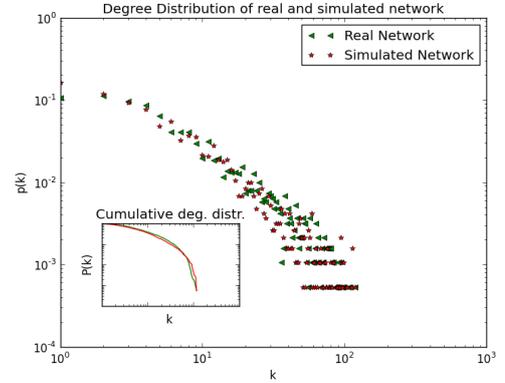


Figure 8: Log-log plot of real (red) and simulated (blue) social network degree distribution $P(k)$ for the Lastfm data set. Inset: plot of the cumulative degree distribution.

We emphasize that Eq.(10) is valuable in predicting users' behavior of a recommender system in an early stage.

5. DISCUSSION

Social influence and our peers are known to form and influence many of our opinions and, ultimately, decisions. We propose here a simple model which is based on heterogeneous agent expectations, a social network, and a formalized social influence mechanism. We analyze the model by numerical simulations and by master equation approach which is particularly suitable to describe the initial phase of the social "contagion". The proposed model is able to generate a wide range of different attendance distributions, including those observed in popular real systems (Netflix, Lastfm, and Movielens). In addition, we showed that these patterns are emergent properties of the dynamics and not imposed by topology of the underlying social network. Of particular interest is the case of Lastfm where the underlying social network is known. Calibrating the observed attendance dis-

tribution against the model then leads not only to social influence parameters but also to the degree distribution of the social network which agrees with that of the true social network.

The Kullback-Leibler distances (KL) for the simulated and real attendance distributions are below 0.05 in all cases, thus demonstrating a good fit. However, the maximum attendances could not be reproduced exactly by the model. One reason may be missing degree correlations in the simulated networks in contrast to real networks where positive degree correlations (so-called degree assortativity) are common. For the Lastfm user friendship network we observe a higher clustering coefficient $C \approx 0.18$ compared to the clustering coefficient $C \approx 0.06$ in the simulated network. To compensate for this, a higher trust parameter γ is needed to fit the real Lastfm attendance distribution with simulated networks.

We are aware that our statistics to validate the model is not complete. But we are confident, that our approach points to a fruitful research direction to understand recommender systems' data as a social driven process.

The proposed model can be a first step towards a data generator to simulate bipartite user-object data with real-world data properties. This could be used to test and validate new recommender algorithms and methods. Future research directions may expand the proposed model to generate ratings within a predefined scale. Moreover, it could be very interesting to investigate the model in the scope of social imitation [36].

6. REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, pages 734–749, 2005.
- [2] R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: an axiomatic approach. In *Proceeding of the 17th international conference on World Wide Web*, pages 199–208. ACM, 2008.
- [3] E. Anderson and L. Salisbury. The formation of market-level expectations and its covariates. *Journal of Consumer Research*, 30(1):115–124, 2003.
- [4] J. Arndt. Role of product-related conversations in the diffusion of a new product. *Journal of Marketing Research*, 4(3):291–295, 1967.
- [5] M. Balabanović and Y. Shoham. Fab: content-based, collaborative recommendation. *Communications of the ACM*, 40(3):66–72, 1997.
- [6] A. Barrat, M. Barthlemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press New York, NY, USA, 2008.
- [7] S. Bergamaschi, F. Guerra, and B. Leiba. Information overload. *Internet Computing, IEEE*, 14(6):10–13, 2010.
- [8] D. Billsus and M. Pazzani. Learning collaborative information filters. In *Proceedings of the Fifteenth International Conference on Machine Learning*, volume 54, page 48, 1998.
- [9] A. Birukov, E. Blanzieri, and P. Giorgini. Implicit: An agent-based recommendation system for web search. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 618–624. ACM, 2005.
- [10] M. Blattner. B-rank: A top N recommendation algorithm. In *Proceedings of The International Multi-Conference on Complexity, Informatics and Cybernetics (IMCIC 2010)*, volume 1, pages 337–341, Orlando, USA, 2010.
- [11] J. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the 14th Annual Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 43–52, San Francisco, CA, 1998. Morgan Kaufmann.
- [12] P. Brusilovsky, A. Kobsa, and W. Nejdl. *The adaptive web: methods and strategies of web personalization*. Springer-Verlag New York Inc, 2007.
- [13] G. Caron-Lormier, R. Humphry, D. Bohan, C. Hawes, and P. Thorbek. Asynchronous and synchronous updating in individual-based models. *ecological modelling*, 212(3-4):522–527, 2008.
- [14] M. Claypool, A. Gokhale, T. Miranda, P. Murnikov, D. Netes, and M. Sartin. Combining content-based and collaborative filters in an online newspaper. In *Proc. ACM SIGIR 99, Workshop Recommender Systems: Algorithms and Evaluation*, 1999.
- [15] H. Drachler, T. Bogers, R. Vuorikari, K. Verbert, E. Duval, N. Manouselis, G. Beham, S. Lindstaedt, H. Stern, M. Friedrich, et al. Issues and considerations regarding sharable data sets for recommender systems in technology enhanced learning. *Procedia Computer Science*, 1(2):2849–2858, 2010.
- [16] F. Fouss, A. Pirotte, J. Renders, and M. Saerens. Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation. *Knowledge and Data Engineering, IEEE Transactions on*, 19(3):355–369, 2007.
- [17] F. Fouss, L. Yen, A. Pirotte, and M. Saerens. An experimental investigation of graph kernels on a collaborative recommendation task. In *Data Mining, 2006. ICDM'06. Sixth International Conference on*, pages 863–868. IEEE, 2007.
- [18] W. Geyer, J. Freyne, B. Mobasher, S. Anand, and C. Dugan. 2nd workshop on recommender systems and the social web. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 379–380. ACM, 2010.
- [19] J. Gleeson. High-accuracy approximation of binary-state dynamics on networks. *Physical Review Letters*, 107(6):68701, 2011.
- [20] D. Goldberg, B. D. Nichols, and D. Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35(12):61–70, 1992.
- [21] N. Good, J. Schafer, J. Konstan, A. Brochers, B. Sarwar, J. Herlocker, and J. Riedl. Combining collaborative filtering with personal agents for better recommendations. In *Proc. Conf. Am. Assoc. Artificial Intelligence (AAAI-99)*, pages 439–446, USA, 1999.
- [22] I. Guy, A. Jaimes, P. Agulló, P. Moore, P. Nandy, C. Nastar, and H. Schinzel. Will recommenders kill

- search?: recommender systems-an industry perspective. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 7–12. ACM, 2010.
- [23] T. Hennig-Thurau, K. Gwinner, G. Walsh, and D. Gremler. Electronic word-of-mouth via consumer-opinion platforms: What motivates consumers to articulate themselves on the Internet? *Journal of Interactive Marketing*, 18(1):38–52, 2004.
- [24] J. Herlocker, J. Konstan, L. Terveen, and J. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, 2004.
- [25] D. Jannach, W. Geyer, J. Freyne, S. Anand, C. Dugan, B. Mobasher, and A. Kobsa. Recommender Systems & the Social Web. In *Proceedings of the 2009 ACM Conference on Recommender Systems, RecSys 2009, New York, NY, USA, October 23-25, 2009*. ACM, 2009.
- [26] K. Goldberg, T. Roeder, D. Gupta, and C. Perkins. Eigentaste: a constant time collaborative filtering algorithm. *Information Retrieval*, (4):133–151, 2001.
- [27] Y. Kim and J. Srivastava. Impact of social influence in e-commerce decision making. *Proceedings of the ninth international conference on Electronic commerce (ICEC)*, pages 293–302, 2007.
- [28] J. Konstan, B. Miller, D. Maltz, J. Herlocker, and L. Gordon. Grouplens: Applying collaborative filtering to usenet news. *Comm. ACM*, 40(3):77–87, 1997.
- [29] S. Kullback and R. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [30] G. Linden, B. Smith, and J. York. Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing*, 7(1):76–80, 2003.
- [31] M. Mason, R. Dyer, and M. Norton. Neural mechanisms of social influence. *Organizational Behavior and Human Decision Processes*, 110(2):152–159, 2009.
- [32] P. Massa and P. Avesani. Trust-aware collaborative filtering for recommender systems. *On the Move to Meaningful Internet Systems 2004: CoopIS, DOA, and ODBASE*, pages 492–508, 2004.
- [33] P. Massa and B. Bhattacharjee. Using trust in recommender systems: an experimental analysis. *Trust Management*, pages 221–235, 2004.
- [34] M.E.J. Newman. Structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [35] P. Melville, R. Mooney, and R. Nagarajan. Content-boosted collaborative filtering for improved recommendations. In *Proc. 18th Nat'l Conf. Artificial Intelligence*, 2002.
- [36] Q. Michard and J. Bouchaud. Theory of collective opinion shifts: from smooth trends to abrupt swings. *The European Physical Journal B-Condensed Matter and Complex Systems*, 47(1):151–159, 2005.
- [37] B. Mirza, B. Keller, and N. Ramakrishnan. Studying recommendation algorithms by graph analysis. *Journal of Intelligent Information Systems*, 20(2):131–160, 2003.
- [38] M. Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.
- [39] J. O'Donovan and B. Smyth. Trust in recommender systems. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 167–174. ACM, 2005.
- [40] M. Pazzani and D. Billsus. Content-based recommendation systems. *Lecture Notes Computer Science*, 4321:325–341, 2007.
- [41] P. Resnick, N. Iakovou, M. Sushak, P. Bergstrom, and J. Riedl. Grouplens: An open architecture for collaborative filtering of networks. In *Proc. Computer Supported Cooperative Work Conf.*, 1994.
- [42] P. Resnick and H. Varian. Recommender systems. *Commun. ACM*, 40:56–58, March 1997.
- [43] F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors. *Recommender Systems Handbook*. Springer, 2011.
- [44] M. Richins and T. Root-Shaffer. THE ROLE OF INVOLVEMENT AND OPINION LEADERSHIP IN CONSUMER WORD-OF-MOUTH: AN IMPLICIT MODEL MADE EXPLICIT. *Advances in consumer research*, 15:32–36, 1988.
- [45] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl. Item-based collaborative filtering recommendation algorithms. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 285–295, New York, NY, USA, 2001. ACM Press.
- [46] K. Sugiyama, K. Hatano, and M. Yoshikawa. Adaptive web search based on user profile constructed without any effort from users. In *Proceedings of the 13th international conference on World Wide Web*, pages 675–684. ACM, 2004.
- [47] F. Walter, S. Battiston, and F. Schweitzer. A model of a trust-based recommendation system on a social network. *Autonomous Agents and Multi-Agent Systems*, 16(1):57–74, 2008.
- [48] G. Webb, M. Pazzani, and D. Billsus. Machine learning for user modeling. *User Modeling and User-Adapted Interaction*, 11(1):19–29, 2001.
- [49] W. Whyte Jr. The web of word of mouth, 1954.
- [50] T. Zhang and V. Iyengar. Recommender systems using linear classifiers. *The Journal of Machine Learning Research*, 2:334, 2002.
- [51] Y. Zhang, M. Blattner, and Y. Yu. Heat conduction process on community networks as a recommendation model. *Physical review letters*, 99(15):154301, 2007.
- [52] T. Zhou, J. Ren, M. Medo, and Y. Zhang. Bipartite network projection and personal recommendation. *Physical Review E*, 76(4):46115, 2007.
- [53] D. Zwillinger and S. Kokoska. *CRC standard probability and statistics tables and formulae*. CRC, 2000.