

Degenerate Feedback Loops in Recommender Systems

Ray Jiang, Silvia Chiappa, Tor Lattimore, Andras Agyorgy, Pushmeet Kohli

Abstract

Machine learning is used extensively in recommender systems deployed in products. The decisions made by these systems can influence user beliefs and preferences which in turn affect the feedback the learning system receives - thus creating a feedback loop. This phenomenon can give rise to the so-called “echo chambers” or “filter bubbles” that have user and societal implications. In this paper, we provide a novel theoretical analysis that examines both the role of user dynamics and the behavior of recommender systems, disentangling the echo chamber from the filter bubble effect. In addition, we offer practical solutions to slow down system degeneracy. Our study contributes toward understanding and developing solutions to commonly cited issues in the complex temporal scenario, an area that is still largely unexplored.

Introduction

Recommender systems are increasingly used to provide users with personalized product and information offerings (Ben Schafer, Konstan, and Riedl, 2001; Lu et al., 2015; Covington, Adams, and Sargin, 2016). These systems employ user’s personal characteristics and past behavior, and similarities among users to generate a list of items that are individually tailored to the user’s preferences. Whilst extremely successful commercially, there are growing concerns that such systems might lead to a self-reinforcing pattern of narrowing exposure and shift in user’s interest, problems that are often referred to in the literature as “echo chamber” and “filter bubble”. A significant amount of recommender systems research has therefore been devoted to deriving ways to favor diversity in the set of items an individual may be exposed to (see Kunaver and Porl (2017) for a review). However, current understanding of the echo chamber and filter bubble effects is limited and experimental analysis reports conflicting results.

In this paper, we provide the first theoretical treatment of these problems. We view user’s interest as a dynamical system and treat interest extremes as degeneracy points of the system. We consider different mod-

els of dynamics and identify sets of sufficient conditions that make them degenerate over time. We then use this analysis to understand the role played by the recommender system. This bipartite approach enables us to separate the phenomenon of the echo chamber effect from the filter bubble effect. Specifically, we generalize the definition of echo chamber in Sunstein (2009), where the term is used to refer to over- and limited-exposure to similar political opinions reinforcing one’s existing beliefs, to refer to user interest being positively or negatively reinforced by repeated exposure to a certain item or category of items. We also focus the definition of filter bubble introduced by Pariser (2011) to describe just the fact that recommender systems select limited content to serve users online.

Finally, we showcase the interplay between the user’s dynamics and the recommender system actions in a simulation study using synthetic data and several classic bandit algorithms. The results reveal several pitfalls of recommender system design and hint towards mitigation strategies.

Related Work

Through an analysis on the MovieLens dataset, Nguyen et al. (2014) found that the diversity of items recommended, and those users engaged with, gets narrower over time. The paper asks whether there is a “natural” tendency of degeneration in user interest. Our paper takes steps toward answering this question by providing theoretical conditions for user interest degeneracy.

In the social sciences literature, Flaxman, Goel, and Rao (2016) found that online services are associated with increased political polarization between users as well as increased exposure to the less preferred side of political opinions. Their seemingly counter-intuitive findings are not contradictory according to our results: systems with some level of random exploration can be degenerative. Barberá et al. (2015) also presented evidence of echo chamber related to political issues on Twitter. On the other hand, Borgesius et al. (2016); Beam, Hutchens, and Hmielowski (2018); Nechush-tai and Lewis (2018) found counter-evidence on online news consumption. Another work by Bakshy, Messing,

and Adamic (2015) measured the effect of user choices separately from that of the recommendation algorithm, and found that individual choices play a larger role than the algorithm in creating echo chamber on Facebook. This supports our viewpoint that user interests degenerate or not depending on their internal dynamics, the recommender system can only slow down or accelerate the process of degeneration.

Model

We consider a recommender system that interacts with a user over time¹. At every time step t , the recommender system serves l items (or categories of items, e.g. news articles, videos, or consumer products) to a user from a finite or countably infinite item² set \mathcal{M} . In general, the goal of the system is to present items to a user that she is interested in: we assume that at time step t , the user’s interest in an item $a \in \mathcal{M}$ is described by a function $\mu_t : \mathcal{M} \rightarrow \mathbb{R}$ such that $\mu_t(a)$ is large (positive) if the user is interested in the item, and it is small (negative) if she is not³.

Given a recommendation $a_t = (a_t^1, \dots, a_t^l) \in \mathcal{M}^l$, the user provides some feedback c_t based on her current interests $\mu_t(a_t^1), \dots, \mu_t(a_t^l)$. This interaction has multiple effects: in the traditional literature for recommender systems, the feedback c_t is used to update the internal model θ_t of the recommender system that has been used to obtain the recommendation a_t , and the new model θ_{t+1} may depend on θ_t , a_t , and c_t . In practice θ_t usually predicts the distribution of user feedback to determine which items a_t should be presented to the user. In this paper we focus on another effect and consider explicitly that the user’s interaction with the recommender system may change her interest in different items for the next interaction, thus the interest μ_{t+1} may depend on μ_t , a_t , and c_t . The full model of interaction is depicted in Fig. 1.

We are interested in studying the evolution of the user’s interest. An example of such an evolution is that the interest is reinforced by user interactions with the recommended items, that is, $\mu_{t+1}(a) > \mu_t(a)$ if the user clicks on an item a at time step t , while $\mu_{t+1}(a) < \mu_t(a)$ if a is shown but not clicked (here $c_t \in \{0, 1\}^l$ can be defined as the indicator vector of clicks to the corresponding items).

To analyze the echo chamber or filter bubble effect, we are interested in understanding when the user’s interest changes extremely, which, in our model, translates to $\mu_t(a)$ taking values arbitrarily different from the

¹For simplicity, in this paper we restrict ourselves to the case of a single user, and leave the case of multiple users possibly influencing each-other interests to future work.

²Throughout the paper, “items” also mean categories of items.

³Whilst we focus on $\mu_t(a) \in \mathbb{R}$, we show in Remark 5, Appendix B that our results can be extended to the case where μ_t belongs to a bounded open interval.

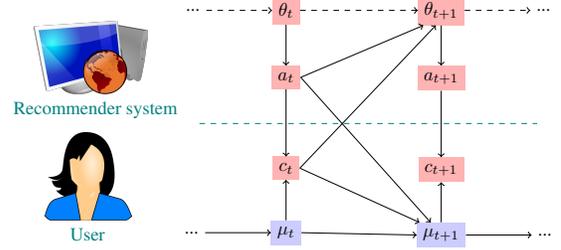


Figure 1: Model of interaction between a recommender system and user over time. Continuous and dashed links indicate existing or possible dependencies, respectively.

initial interest $\mu_0(a)$: large positive values indicate that the user becomes extremely interested in item a , while large negative values indicate that the user dislikes a . Formally, for a finite \mathcal{M} , we can ask if the L^2 norm $\|\mu_t - \mu_0\|_2 = (\sum_{a \in \mathcal{M}} (\mu_t(a) - \mu_0(a))^2)^{1/2}$ can grow arbitrarily large: the user’s interest sequence μ_t is called *weakly degenerate* if

$$\limsup_{t \rightarrow \infty} \|\mu_t - \mu_0\|_2 = \infty \quad \text{almost surely}^4. \quad (1)$$

We show in the next section that weak degeneracy occurs under mild conditions on the dynamics of μ_t (Theorem 1). A stronger notion of degeneracy, which also requires that once μ_t drifted away from μ_0 it remains so, is *strong degeneracy*: the sequence μ_t is *strongly degenerate* if

$$\lim_{t \rightarrow \infty} \|\mu_t - \mu_0\|_2 = \infty \quad \text{almost surely}. \quad (2)$$

In the next section we show that under mild assumptions on the evolution of the user’s interest sequence μ_t becomes degenerate.

There are multiple ways to extend the above definitions to the case of an infinite item set \mathcal{M} . For simplicity, we only consider here replacing $\|\mu_t - \mu_0\|_2$ with $\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)|$ in Eqs. (1) and (2), which is equivalent to the original definitions when \mathcal{M} is finite⁵.

User Interest Dynamics – Echo Chamber

As items often represent diverse categories of things, we make the simplifying assumption that they are independent from each other. By setting $l = 1$ and $a_t^1 = a$ for all t (i.e., $\mathcal{M} = \{a\}$), we can remove the influence of the recommender system and consider the dynamics of the user’s interest separately. This allows us to analyze the echo chamber effect: what happens to the interest $\mu_t(a)$ if item a is served infinitely often (i.o.).

Since a is fixed, to simplify the notation, we write μ_t instead of $\mu_t(a)$ in this section. Given a_t , according

⁴i.e. with probably 1.

⁵As such, we could have used $\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)|$ in our original definitions, but we prefer $\|\mu_t - \mu_0\|_2$ as it also provides some information about the “average” deviation of the user’s interest over the different items at any finite time t .

to Fig. 1, μ_{t+1} is a—possibly stochastic—function of μ_t (as μ_{t+1} depends on c_t and μ_t , and c_t depends on μ_t). Below we discuss the general case when the drift $\mu_{t+1} - \mu_t$ is a nonlinear stochastic function; deterministic models for the drift are considered in Appendix B.

Nonlinear Stochastic Model. We assume that $\mu_0 \in \mathbb{R}$ is fixed and that $\mu_{t+1} = \mu_t + f(\mu_t, \xi_t)$, where $(\xi_t)_{t=1}^\infty$ is an infinite sequence of independent uniformly distributed random variables that introduce noise into the system (*i.e.* μ_{t+1} is a stochastic function of μ_t). The function $f : \mathbb{R} \times [0, 1]$ is assumed to be measurable, but otherwise arbitrary. Denoting the uniform distribution over $[0, 1]$ by $U([0, 1])$, let

$$\bar{f}(\mu) = \mathbb{E}_{\xi \sim U([0,1])}[f(\mu, \xi)]$$

be the expected increment $\mu_{t+1} - \mu_t$ when $\mu_t = \mu$. We also define

$$F(\mu, x) = \mathbb{P}_{\xi \sim U([0,1])}(f(\mu, \xi) \leq x)$$

to be the cumulative distribution of the increment. The asymptotic behavior of μ_t depends on f , but under mild assumptions the system degenerates weakly (Theorem 1) or strongly (Theorem 2)⁶.

Theorem 1 (weak degeneracy). *Assume that F is continuous at $(\mu, 0)$ for all $\mu \in \mathbb{R}$ and there exists a $\mu_\circ \in \mathbb{R}$ such that 1) $F(\mu, 0) < 1$ for all $\mu \geq \mu_\circ$, 2) $F(\mu, 0) > 0$ for all $\mu \leq \mu_\circ$. Then the sequence μ_t is weakly degenerate, that is, $\limsup_{t \rightarrow \infty} |\mu_t| = \infty$ almost surely.*

The assumptions guarantee that within any closed bounded interval there is a constant probability that the random walk escapes to the left/right when starting to the left/right of μ_\circ respectively. Under stronger conditions it is possible to guarantee the divergence of the random walk. We state a simple version of the theorem, but note that the result can be generalized in many ways.

Theorem 2 (strong degeneracy). *Assume that the conditions of Theorem 1 hold, and additionally that there exists $c \in \mathbb{R}$ such that $|\mu_{t+1} - \mu_t| \leq c$ almost surely and there exists an $\epsilon > 0$ such that for all sufficiently large μ it holds that $\bar{f}(\mu) > \epsilon$, and for all sufficiently small μ it holds that $\bar{f}(\mu) < -\epsilon$. Then $\lim_{t \rightarrow \infty} \mu_t = \infty$ or $\lim_{t \rightarrow \infty} \mu_t = -\infty$ almost surely.*

Intuitively, weak degeneration occurs in a stochastic environment if the user’s interest has some non-zero probability of drifting up when above some threshold, and of drifting down when below. Strong degeneracy holds if additionally $|\mu_{t+1} - \mu_t|$ is bounded and for μ_t sufficiently large/small the increment $\mu_{t+1} - \mu_t$ has positive/negative drift that is larger than a constant.

Theorems 1 and 2 show that the user’s interest degenerates under very mild conditions, in particular, in the model we consider in our simulation studies. Thus, in such cases degeneracy can only be avoided if an item (or item category) is showed only finitely many times;

⁶The proofs of these theorems are given in Appendix A.

otherwise one can only hope to control how fast μ_t degenerates (*i.e.* tends to ∞).

System Design Role – Filter Bubble

In the previous section we discussed conditions for degeneracy for different user interest dynamics. In this section we examine the other side of the story, the influence of recommender system actions in creating filter bubbles. We typically do not know the dynamics of the user’s interest in the real world. However, we consider the relevant scenario to the echo chamber/filter bubble problem where user’s interest in some items has degenerate dynamics, and examine how to design a recommender system that slows down the degeneracy process. We consider three dimensions, namely model accuracy, amount of exploration, and growing candidate pool.

Model Accuracy. One common goal of recommender systems designers is to increase the prediction accuracy of the internal model θ_t . How does model accuracy coupled with greedy optimal a_t affect the speed of degeneration? We examine this question for the extreme case of exact predictions, *i.e.* $\theta_t = \mu_t$, we call such a prediction model the *oracle model*. We argue that under the *surfacing assumption* explained below, the oracle model coupled with greedily optimal action selection is the quickest path to degeneracy.

In order to analyze the problem concretely, we focus on the degenerate linear dynamics model for $\mu_t(a)$ for $a \in \mathcal{M}$, *i.e.* $\mu_{t+1}(a) = (1 + k)\mu_t(a) + b$. Then we can solve for $\mu_t(a)$ and $\mu_t(a) = (\mu_0(a) + \frac{b(a)}{k(a)})(1 + k(a))^t - \frac{b(a)}{k(a)}$, for $|1 + k(a)| > 1$ (see Appendix B).

Surfacing Assumption: Let $[m] = \{1, 2, \dots, m\}$ be the candidate set of size m . If a subset of items $\mathcal{S} \subset [m]$ leads to positive degenerate dynamics (*i.e.* $\mu_t(a) \rightarrow +\infty$ for all $a \in \mathcal{S}$), then we assume that there exists a time $\tau > 0$ such that, for all $t \geq \tau$, \mathcal{S} takes up the top $|\mathcal{S}|$ items in terms of values of μ_t , sorted by the base value of the exponential function, $|1 + k(a)|$.

The surfacing assumption makes sure that the quickest degenerating items surface out to the top list given enough time of exposure. It can be generalized to nonlinear stochastic dynamics of μ_t provided that the items from \mathcal{S} have an almost surely stable ordering of degeneracy speed $|\mu_t(a) - \mu_0(a)|/t$ over time.

Under the general surfacing assumption, after time τ , the quickest way to degeneration is to serve the top l items according to μ_t , or θ_t of the oracle model. Even if the assumption is violated to some degree, the oracle model still leads to degeneracy very efficiently by picking the top l items according to μ_t which are likely to receive positive feedback due to high μ_t , and therefore increasing μ_{t+1} and reinforcing the past choices.

In practice the recommender system models are inaccurate. We can think of inaccurate models as the oracle model with different levels of noises added to θ_t . Hence we discuss inaccurate models in the next section.

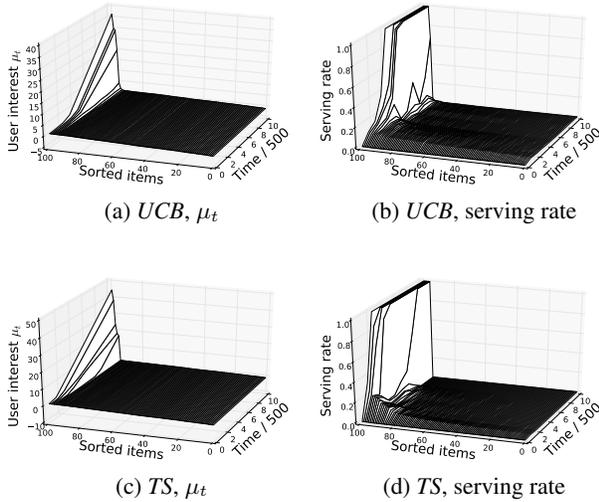


Figure 2: Echo chamber and filter bubble effect for UCB and TS models. Sorted user interest μ_t and serving rates are plotted every 500 steps.

Amount of Exploration. Consider a type of ϵ -random exploration where a_t always picks the top l items out of a finite candidate pool $[m]$ with uniform ϵ noise on θ_t , *i.e.* according to $\theta_t^l = \theta_t + U([- \epsilon, \epsilon])$.

Given the same model sequence θ_t , the bigger ϵ is, usually the slower the system degenerates. However, in practice θ_t is learned from observations, and the random exploration added to an oracle model may in fact accelerate degeneration: random exploration can help reveal the most positively degenerating items over time making the surfacing assumption more likely to be true (we show this phenomenon in the simulation experiments below, Fig. 5). In addition, if user interests have degenerative dynamics, even recommending items uniformly at random leads to degeneration, albeit quite slowly.

How do we then make sure that the recommender system does not degenerate? One way is to limit the number of times an item for which the user’s interest dynamics is degenerative is served to the user. In practice it is hard to detect which items correspond to degenerative dynamics, however we can generally prevent degeneration if all items are served only a finite number of times, which suggests having an ever growing pool of candidate items.

Growing Candidate Pool \mathcal{M} . With a growing candidate pool, at every time step, an additional set of new items become available to be served to user. Hence the domain of the function μ_t expands as t increases. Adding new items at least linearly often is a necessary condition to escape system degeneration, since in a finite or any sublinearly growing candidate pool, by the pigeon hole principle, there must exist at least one item that is served *i.o.*, which is degenerate in the worst case scenario (also under general conditions described *e.g.* in

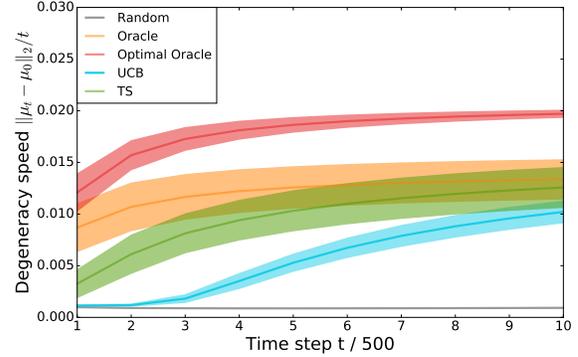


Figure 3: System evolves for 5,000 time steps with report interval 500. The results are averaged over 30 runs with the shaded area indicating the standard deviation.

Theorem 2). However, with an at least linearly growing candidate pool \mathcal{M} the system can potentially impose the maximum number of times any item is served to a user and prevent the system from degeneration.

Simulation Experiments

In this section, we create a simple degenerate dynamics of μ_t and examine the degeneration speed under five different recommender system models. We further demonstrate that adding new items to the candidate pool can be an effective solution against system degeneration.

We create a simulation for the model of interaction between a recommender system and a user of Fig. 1. Consider a possibly growing candidate pool of items of initial size m_0 and of size m_t at time step t . At each time step t , a recommender system picks the top l out of the m_t items $a_t = (a_t^1, \dots, a_t^l)$ according to the internal model θ_t to serve to a user. The user considers each of the l items independently and chooses to click on a (possibly empty) subset of them, thereby generating a binary vector c_t of size l where $c_t(a_t^i)$ gives the user feedback on item a_t^i , according to $c_t(a_t^i) \sim \text{Bernoulli}(\phi(\mu_t(a_t^i)))$, where ϕ is the sigmoid function $\phi(x) = 1/(1 + e^{-x})$. The system then updates the model θ_{t+1} based on the past actions, feedbacks and the current model parameter θ_t . We consider one of the simplest degenerate dynamics, where the user’s interest is modeled by a constant increase $\delta(a')$ if the item a' receives a click and a constant decrease otherwise. The function δ maps from the candidate set to \mathbb{R} :

$$\mu_{t+1}(a_t^i) - \mu_t(a_t^i) = \begin{cases} \delta(a_t^i) & \text{if } c_t(a_t^i) = 1, \\ -\delta(a_t^i) & \text{otherwise.} \end{cases} \quad (3)$$

From Theorem 2, we know that $\mu_t \rightarrow \pm\infty$ for every item. In the experiment, we set $l = 5$ and sample the drift δ from a uniform random distribution $U([-0.01, 0.01])$. The user’s initial interest μ_0 for all items is independently sampled from a uniform random distribution $U([-1, 1])$.

The internal recommender system model is updated

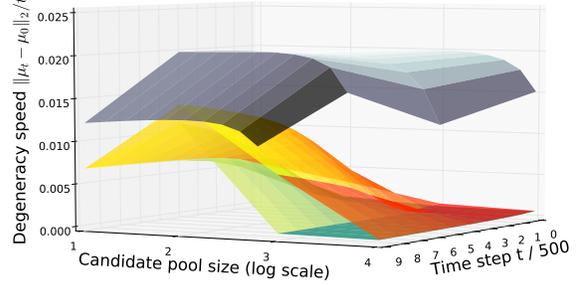
according to following five algorithms:

- **Random Model:** Instead of picking top items, the set of items $a_t = (a_t^1, \dots, a_t^l)$ is sampled from a uniform random distribution over the candidate set $U(\{m_t\})$.
- **Oracle:** $\theta_{t+1}(a_t^i) = \mu_{t+1}(a_t^i), \forall i$.
- **Optimal Oracle:** $\theta_{t+1}(a_t^i) = \delta(a_t^i), \forall i$. This model does not pick the highest l items according to μ_t but according to δ . Thus, it always picks the fastest degenerating items, therefore maximizing both the long term user engagement $\sum_{t=0}^{\infty} \|c_t\|_1$ and the degeneracy speed. For a fixed candidate pool \mathcal{M} , this model is equivalent to an *Oracle* that satisfies the Surfacing Assumption.
- **Upper Confidence Bound Multi-armed Bandit Algorithm (UCB)** (Lai, 1987; Auer, Cesa-Bianchi, and Fischer, 2002; Lattimore and Szepesvári, 2019): We use the version of UCB algorithm in Chapter 8 of Lattimore and Szepesvári (2019), however most UCB algorithms perform similarly to the purpose of this experiment. The algorithm prioritizes serving any item from the candidate set that has never been served before. This treatment includes the initial m_0 items as well as later whenever new items are added to the candidate pool. At time step t , UCB serves $l'(0 \leq l' \leq l)$ previously unserved items and the top $l - l'$ items according to values of θ_t . Define $f(t) = 1 + t \log^2(t)$ and we use the following model update $\theta_{t+1}(a) = \hat{c}_t(a) + \sqrt{2 \log f(t)/T_a(t)}$, where \hat{c}_t is the empirical average of feedbacks on item a , i.e. $\hat{c}_t(a) = \sum_{0 \leq i \leq t, a \in a_i} c_t(a)/T_a(t)$, and $T_a(t)$ is the number of times item a has been served up to time t , i.e. $T_a(t) = \sum_{0 \leq i \leq t, a \in a_i} 1$.
- **Thompson Sampling Multi-armed Bandit Algorithm (TS)** (Thompson, 1933): We initialize $\alpha_0(a) = 1, \beta_0(a) = 1$ for any new item a . If a is served at time t , we perform the update $\alpha_{t+1}(a) = \alpha_t(a) + c_t(a), \beta_{t+1}(a) = \beta_t(a) + 1 - c_t(a)$. At any time t , the internal model θ_t is sampled from the corresponding beta distribution $Beta(\alpha_t, \beta_t)$.

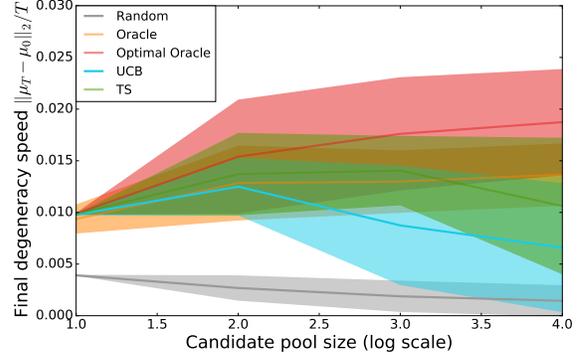
Echo Chamber & Filter Bubble Effect

We examine the echo chamber and filter bubble effects by running the simulation on a candidate pool of fixed size $m_t = m = 100$ with time horizon $T = 5,000$.

In Fig. 2, we show the degeneration of user interest μ_t (left) and the serving rate (right) of every item as *UCB* and *TS* evolves in time (due to space restrictions, the plots for the remaining models are shown in Fig. 7 of Appendix D). The serving rate of an item shows how often it is served within the report interval. In order to see the distribution clearly, we sort the items according to the z-values at the report time. Under both *UCB* and *TS* models, very quickly both the top items served and the top user interests narrow down to the ($l =$) 5 most positively reinforced items.



(a) Degeneracy rate vs candidate pool size vs time



(b) Degeneracy rate vs candidate pool size

Figure 4: (a) Degeneracy surfaces for *Optimal Oracle* (grey), *UCB* (green) and *TS* (orange) up to time $T = 5,000$ while varying candidate pool sizes $\log_{10} m = 1, 2, 3, 4$. (b) Degeneracy speed at $T = 20,000$.

Speed of Degeneracy

Next, we compare the system degeneracy speeds for the five recommender system models on both fixed and growing candidate sets. As the L^2 distance that measures system degeneracy is asymptotically linear for all five models (see Appendix C), we quantify degeneracy speeds by compare empirically $\|\mu_t - \mu_0\|_2/t$ in finite candidate pools for different experiment setups.

Figure 3 shows the degeneracy speed of five models averaged across 30 runs when we take $m = 100$ and evolve the system for $T = 5,000$ steps. We see that the *Optimal Oracle* degenerates the fastest by far, followed by the *Oracle*, *TS* and *UCB*. The *Random Model* offers the slowest degeneracy speed.

The Effect of Candidate Pool Size. In Fig. 4a we compare the *Optimal Oracle*, *UCB* and *TS*' degeneracy speed $\|\mu_t - \mu_0\|_2/t$ up to 5,000 time steps and candidate pool sizes $m = 10, 10^2, 10^3, 10^4$. Apart from the *Random* model, we see that *UCB* slows down system degeneracy the most given a large candidate pool since it is forced to explore any unserved item first. A larger candidate pool requires a longer time for exploration for the bandit algorithms. As the candidate pool size grows to 10,000 *UCB*'s degeneracy speed never peaks up given the time horizon, but will eventually grow

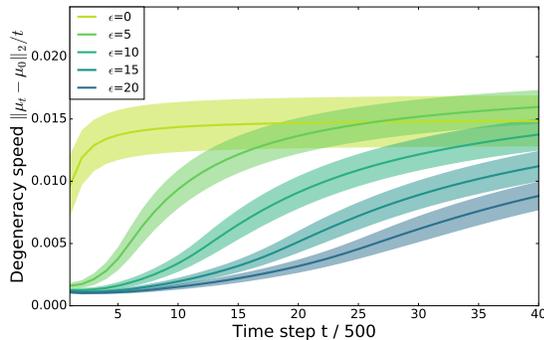


Figure 5: Degeneracy speed for the *Oracle* model with different noise levels $\epsilon \in [0, 10]$ up to $T = 20,000$.

given a longer time. *TS* has higher degeneracy speed due to weaker exploration on new items. The *Optimal Oracle* accelerates degeneration given a larger pool, as it can pick potentially faster degenerative items than from a smaller pool.

Additionally, in Fig. 4b we plot all five models degeneracy speed for $T = 20,000$ against the same changing candidate pool sizes. The degeneracy speed of the *Optimal Oracle* and the *Oracle* increases with the size of the candidate set, but that of the *and Random Model*, *UCB*, and *TS* decreases. In practice, having a large candidate pool can be a temporary solution to slow down system degeneration.

The Effect of the Noise Level. Next we show the influence of internal model inaccuracy on degeneracy speed. We compare the *Oracle* model with different amounts of uniformly random noises, *i.e.* the system serves the top l items according to the noisy internal model $\theta'_t = \theta_t + U([- \epsilon, \epsilon])$. The candidate pool has fixed size $m = 100$. In Fig. 5, we vary ϵ from 0 to 10. Counter-intuitively adding noise to *Oracle* accelerates degeneration since faster degenerative items may be selected by chance than those fixed set of top l items ranked by μ_0 , and more likely satisfies the Surfacing Assumption. Given $\epsilon > 0$, as expected, we see a nice monotonically increasing damping effect on degeneracy speed as the noise level grows.

Growing Candidate Pool. We extend the definition of degeneracy speed to an infinite candidate pool by computing $\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)|/t$ (see Appendix C for an asymptotic analysis). Since the degeneracy speed may not be asymptotically linear for all five models, we examine directly the sup distance $\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)|$ over 10,000 time steps. To construct growing candidate pools at different growth speed, we define a growth function $m_t = \lfloor m_0 + lt^\eta \rfloor$ by varying the growth parameter⁷ $\eta = 0, 0.5, 1$, where $m_0 = 100$. In Fig. 6 we average the results over 10 independent runs. Both the *Oracle* and the *Optimal Oracle* for all growth

⁷ $\eta = 0$ gives a fixed candidate pool, $0 < \eta < 1$ gives sub-linear growth, $\eta = 1$ gives linear growth.

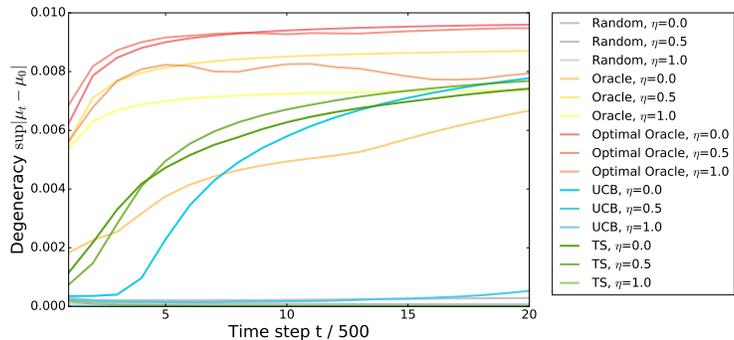


Figure 6: Comparison of the five models with growing candidate pools at different rates $\eta = 0, 0.5, 1.0$, degeneracy up to $T = 10,000$, averaged over 10 runs.

rates are degenerate. The *Random Model* stops degeneration at sublinear growth, $\eta = 0.5$, so does *UCB* thanks to forced exploration on previously unserved items, although its trajectory has a small upward tilt. The *TS* model degenerates at sublinear growth but stops degeneration at linear growth $\eta = 1$. For all models, the higher the growth rate η , the slower they degenerate, if they do at all. Overall when applicable, an ideally linearly growing candidate set and continuous random exploration seem to be good remedies against an adversarial dynamics of μ_t to best prevent degeneration.

Conclusion

We provided a theoretical analysis of the echo chamber and filter bubble effects for recommender systems. We used the dynamical system framework to model user’s interest and treated interest extremes as degeneracy points of the system. We gave formal definitions of system degeneracy and provided sufficient conditions which make the system degenerate with both deterministic and stochastic dynamics. On the recommender system side, we discussed the influence on degeneracy speed of three independent factors in system design, *i.e.* model accuracy, amount of exploration, and the growth rate of the candidate pool. An oracle model often leads to quick degeneracy of the system, while continuous exploration and a large candidate pool size can help slow it down. The best remedies against system degeneracy we found are continuous random exploration and growing the candidate pool at least linearly.

There are two limitations of our work. First, on the MovieLens dataset, on which Nguyen et al. (2014) found evidence of degeneration, our theoretical models gave inconclusive results dependent on the proxy measure used for user interests. A good measure or proxy for user interests is necessary in practice to detect degeneration. Second, we assumed that items and users are independent from each other – we will extend the theoretical analysis to the case of possibly mutually dependent items and users in a future work.

Acknowledgments

We would like to thank William Isaac, Michael Mathieu, Krishnamurthy Dvijotham, Timothy Mann, Dilan Gorur, Ulrich Paquet and Sumedh Ghaisas, for helpful discussions and advice.

References

- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47(2):235–256.
- Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science* 348:1130–1132.
- Barberá, P.; Jost, J. T.; Nagler, J.; Tucker, J. A.; and Bonneau, R. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science* 26(10):1531–1542.
- Beam, M. A.; Hutchens, M. J.; and Hmielowski, J. D. 2018. Facebook news and (de)polarization: reinforcing spirals in the 2016 us election. *Information, Communication & Society* 21(7):940–958.
- Ben Schafer, J.; Konstan, J.; and Riedl, J. 2001. E-commerce recommendation applications. *Data Mining and Knowledge Discovery* 115–153.
- Borgesius, F. J. Z.; Trilling, D.; Möller, J.; Bodó, B.; de Vreese, C. H.; and Helberger, N. 2016. Should we worry about filter bubbles? *Internet Policy Review*.
- Covington, P.; Adams, J.; and Sargin, E. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys)*.
- Flaxman, S.; Goel, S.; and Rao, J. 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80:298–320.
- Galor, O. 2007. *Discrete Dynamical Systems*. Springer.
- Kunaver, M., and Porl, T. 2017. Diversity in recommender systems a survey. *Knowledge-Based Systems* 123(C):154–162.
- Lai, T. L. 1987. Adaptive treatment allocation and the multi-armed bandit problem. *Ann. Statist.* 15(3):1091–1114.
- Lattimore, T., and Szepesvári, C. 2019. *Bandit Algorithms*. Cambridge (to appear).
- Lu, J.; Wu, D.; Mao, M.; Wang, W.; and Zhang, G. 2015. Recommender system application developments: A survey. 74.
- Nechushtai, E., and Lewis, S. C. 2018. What kind of news gatekeepers do we want machines to be? filter bubbles, fragmentation, and the normative dimensions of algorithmic recommendations. *Computers in Human Behavior*.

Nguyen, T.; Hui, P.; Harper, F.; Terveen, L.; and Konstan, J. 2014. Exploring the filter bubble: The effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, 677–686.

Pariser, E. 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin UK.

Sunstein, C. R. 2009. *Republic.com 2.0*. Princeton University Press.

Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25(3/4):285–294.

A Proofs

Proof of Theorem 1. Let \mathbb{P}_μ denote the measure carrying $(\mu_t)_{t=1}^\infty$ when $\mu_0 = \mu$. Assume without loss of generality that $\mu_\circ = 0$ and let $B > 0$ be arbitrary. The result will follow by showing that for all $\mu \in [-B, B]$ it holds that

$$\mathbb{P}_\mu(\text{exists } t : \mu_t \notin [-B, B]) = 1. \quad (4)$$

To see this suppose that

$$\mathbb{P}_\mu(\limsup_{t \rightarrow \infty} |\mu_t| < \infty) > 0.$$

Then there exists a $B > 0$ such that

$$\mathbb{P}_\mu(\text{exists } t : \mu_t \notin [-B, B]) < 1,$$

which is a contradiction. In order to prove (4) notice that compactness of $[0, B]$ and $[-B, 0]$ and the continuity of F at $(\mu, 0)$ for all μ ensures there exists an $\epsilon > 0$ depending only on B such that $1 - F(\mu, \epsilon) \geq \epsilon$ for all $\mu \in [0, B]$ and $F(\mu, -\epsilon) \geq \epsilon$ for all $\mu \in [-B, 0]$. Hence for $n = 1 + B/\epsilon$

$$\inf_{\mu \in [-B, B]} \mathbb{P}_\mu(\text{exists } t \leq n : \mu_t \notin [-B, B]) \geq \epsilon^n.$$

Let E_k be the event that $\mu_t \in [-B, B]$ for all $t \in \{nk + 1, \dots, n(k+1) + 1\}$. Then

$$\begin{aligned} \mathbb{P}_\mu(\mu_t \in [-B, B] \text{ for all } t) &= \mathbb{P}_\mu(E_k \text{ for all } k \geq 0) \\ &= \prod_{k=0}^{\infty} \mathbb{P}_\mu(E_k \mid E_1, \dots, E_{k-1}) \\ &\leq \prod_{k=0}^{\infty} (1 - \epsilon^n) = 0, \end{aligned}$$

which completes the proof. \square

Proof of Theorem 2. We prove that if B is sufficiently large and $\mu > B$ and $\tau = \min\{t : \mu_t \in [-B, B]\}$, then

$$\mathbb{P}_\mu\left(\lim_{t \rightarrow \infty} \mu_t = \infty\right) > \frac{1}{2} \quad (5)$$

and

$$\mathbb{P}_\mu\left(\lim_{t \rightarrow \infty} \mu_t = \infty \text{ or } \tau < \infty\right) = 1. \quad (6)$$

For $\mu < -B$ the same holds, but with μ_t tending to $-\infty$. To see why this implies the result, notice that the previous theorem shows that μ_t eventually leaves $[-B, B]$ almost surely. Each time this happens there is more than 0.5 probability of divergence and a certainty of either divergence or returning to $[-B, B]$. The conditional Borel-Cantelli theorem concludes the proof. To see why (5) and (6) hold, let

$$M_n = \sum_{t=1}^n (\mu_{t+1} - \mu_t - \bar{f}(\mu_t)),$$

which is a martingale with bounded increments since $|\mu_{t+1} - \mu_t|$ are bounded. Now $M_{\tau \wedge n}$ is also a martingale. Then by the strong law of large numbers for martingales,

$$\lim_{n \rightarrow \infty} M_{\tau \wedge n} / n = 0 \text{ a.s.},$$

which implies that either $\tau < \infty$ or $\mu_n \rightarrow \infty$ almost surely. To see the latter case, suppose μ is sufficiently large, we compute

$$\begin{aligned} \mu_{n+1} &= M_n + \mu_1 + \sum_{t=1}^n f(\mu_t) \\ &> M_n + \mu_1 + \epsilon n \rightarrow \infty \end{aligned}$$

since $M_n/n \rightarrow 0$ a.s. A similar argument applies when μ is sufficiently small. Finally, Eq. (5) holds by Azuma's inequality. \square

B More on User Interest Dynamics

Linear Deterministic Model

Recall that since a is fixed, to simplify the notation, we write μ_t, c_t instead of $\mu_t(a), c_t(a)$ in this section. We assume $c_t = g(\mu_t)$ for some deterministic function g and analyze the one-dimensional, autonomous, first-order discrete dynamical system:

$$\mu_{t+1} = \mu_t + h(c_t) = \mu_t + f(\mu_t), \quad (7)$$

where $f = h \circ g$. In the simple case in which f is a linear function, i.e. $f(\mu_t) = k\mu_t + b$ for a linear coefficient k and constant term b , Eq. 7 takes the form

$$\mu_{t+1} = (1+k)\mu_t + b. \quad (8)$$

By unrolling Eq. (8) over time we obtain

$$\mu_t = (1+k)^t \mu_0 + b \sum_{i=0}^{t-1} (1+k)^i \quad (9)$$

$$= \begin{cases} \mu_0 + bt, & \text{if } k = 0 \\ (\mu_0 + \frac{b}{k})(1+k)^t - \frac{b}{k}, & \text{if } k \neq 0. \end{cases} \quad (10)$$

For $k = 0$, the user's interest is not influenced by the recommender system but drifts with the constant b – this case occurs in the real world with probability 0.

For $k \neq 0$, the steady state equilibrium, obtained from solving the equation $\bar{\mu} = \bar{\mu} + k\bar{\mu} + b$, is given by $\bar{\mu} = -\frac{b}{k}$. Substituting $-\frac{b}{k}$ with $\bar{\mu}$ in Eq. 10 and taking the

limit $t \rightarrow \infty$ on both sides, we obtain

$$\lim_{t \rightarrow \infty} \mu_t = (\mu_0 - \bar{\mu}) \lim_{t \rightarrow \infty} (1+k)^t + \bar{\mu}, \quad (11)$$

namely,

$$\lim_{t \rightarrow \infty} \mu_t = \begin{cases} \bar{\mu}, & \text{if } |1+k| < 1, \\ \bar{\mu}, & \text{if } |1+k| > 1 \text{ and } \mu_0 = \bar{\mu}, \\ \text{alternating } \{\mu_0, 2\bar{\mu} - \mu_0\}, & \text{if } k = -2, \\ \infty & \text{if } 1+k > 1, \end{cases} \quad (12)$$

and, $\limsup_{t \rightarrow \infty} \mu_t = \infty$ when $1+k < -1$. (13)

Therefore, no matter how the system selects items, in the first three cases of Eq. (12) the user interest model over the item a in question is always bounded over time. Of these, only the first case $|1+k| < 1$, or equivalently $-2 < k < 0$, will occur with probability different from 0.

On the other hand, for $k > 0$ the recommender system will degenerate strongly with μ_t growing at an exponential rate. For $k < -2$, the recommender system will degenerate weakly with $\sup_{i < t} \mu_i$ growing exponentially.

In summary, we can draw the following conclusions when an item a is served *i.o.*: 1) for $k > 0$ or $k < -2$, the user interest μ_t degenerates by growing exponentially, and therefore the recommender system needs to exert control on how frequent such items are shown to the user in order to control the speed of degeneracy of μ_t ; 2) in all other cases, the user interest μ_t does not degenerate ($-2 \leq k < 0$), or the system cannot control linear degeneracy ($k = 0$, improbable case).

Non-linear Deterministic Model

If f is a non-linear deterministic function, the steady state equilibrium is reached at zeros of f . Sufficient (but not necessary) conditions for global stability are given by the following theorem (Galor, 2007):

Theorem 3 (sufficiency). *If $g : \mathbb{R} \rightarrow \mathbb{R}$ is a contraction mapping, i.e. if*

$$\frac{|g(y_{t+1}) - g(y_t)|}{|y_{t+1} - y_t|} < 1, \quad \forall t = 0, 1, 2, \dots, \infty,$$

then a stationary equilibrium of the difference equation $y_{t+1} = g(y_t)$ exists and is unique and globally (asymptotically) stable.

Thus if $g = I + f$ (where I is the identity function) is a contraction mapping, then there exists a steady state equilibrium that is unique and globally stable, and the system will not degenerate. Both requiring a globally stable steady state equilibrium and requiring g to be a contraction mapping are strong sufficient conditions. Moreover it is almost always impossible to verify it in practice since we don't know the actual function g . Thus we gave three examples of sufficient conditions that are used to describe the general dynamics of user interests. In the first example, users respond to an item if

their interests level exceed some action threshold. With each feedback, their interest level changes and so does the action threshold. If after a while the changes in action thresholds will always be smaller than user interest changes and the user interest's total variation ranges over \mathbb{R} , then the system degenerates.

Theorem 4 (sufficiency). *Let $\{d_1, d_2, \dots\}$ be an infinite sequence, $d_t \in \mathbb{R}$. Then if the user interest for an item has the following dynamics,*

$$\mu_t > d_t \iff \mu_{t+1} > \mu_t, \quad (14)$$

$$\exists t_0 > 0 \text{ s.t. } |d_{t+1} - d_t| \leq |\mu_{t+1} - \mu_t|, \forall t \geq t_0, \quad (15)$$

$$\sum_{t=0}^{\infty} |\mu_{t+1} - \mu_t| = \infty, \quad (16)$$

then it degenerates strongly as $t \rightarrow \infty$.

Proof. We prove that $\mu_t \rightarrow \pm\infty$ as $t \rightarrow \infty$. At time t_0 , either $\mu_{t_0} > d_{t_0}$ or $\mu_{t_0} \leq d_{t_0}$. First we consider the case where $\mu_{t_0} > d_{t_0}$. At the next time step $t_0 + 1$, there are again two different cases:

1. $d_{t_0+1} < d_{t_0}$: it implies that $\mu_{t_0+1} > \mu_{t_0} > d_{t_0} > d_{t_0+1}$ by Condition 14.
2. $d_{t_0+1} \geq d_{t_0}$: by Condition 15, $d_{t_0+1} - d_{t_0} \leq \mu_{t_0+1} - \mu_{t_0}$ and thus

$$\mu_{t_0+1} \geq d_{t_0+1} - d_{t_0} + \mu_{t_0} > d_{t_0+1}$$

using $\mu_{t_0} - d_{t_0} > 0$. Hence in both cases we have $\mu_{t_0+1} > d_{t_0+1}$. Applying the same argument to every time step, we also have

$$\begin{aligned} \mu_{t_0} > d_{t_0} &\implies \mu_{t_0+1} > d_{t_0+1} \\ &\implies \mu_{t_0+2} > d_{t_0+2} \\ &\implies \dots \end{aligned}$$

which implies

$$\mu_{t_0} < \mu_{t_0+1} < \mu_{t_0+2} < \dots$$

By Condition 16, we conclude that if $\mu_{t_0} > d_{t_0}$ then $\mu_t \rightarrow \infty$ as $t \rightarrow \infty$.

In the other case, $\mu_{t_0} \leq d_{t_0} \iff \mu_{t_0+1} \leq \mu_{t_0}$. Similarly at time $t_0 + 1$, we have either $d_{t_0+1} < d_{t_0}$ or $d_{t_0+1} \geq d_{t_0}$. Following the same argument as above and reversing the inequality signs, we have

$$\begin{aligned} \mu_{t_0} \leq d_{t_0} &\implies \mu_{t_0+1} \leq d_{t_0+1} \\ &\implies \mu_{t_0+2} \leq d_{t_0+2} \\ &\implies \dots \end{aligned}$$

which implies

$$\mu_{t_0} \leq \mu_{t_0+1} \leq \mu_{t_0+2} \leq \dots$$

By Condition 16, we conclude that if $\mu_{t_0} \leq d_{t_0}$ then $\mu_t \rightarrow -\infty$ as $t \rightarrow \infty$. \square

Scale-invariance

Let \mathcal{M} be the candidate pool of items, which is also the domain of function μ_t .

Remark 5 (scale-invariance). *Let $\mu_t : \mathcal{M} \rightarrow D$ for any open interval $D \subset \mathbb{R}$. All sufficiency theorems still hold if we define the system degeneracy for μ_t as the following. We call $\mu_t : \mathcal{M} \rightarrow D$ degenerate iff there exists a monotonic, continuous function $\psi : D \rightarrow \mathbb{R}$ such that $\limsup_{t \rightarrow \infty} \|\psi \circ \mu_t - \psi \circ \mu_0\|_2 = \infty$.*

If we define $\nu_t = \psi \circ \mu_t$, then all sufficiency theorems readily apply to ν_t and the conclusions hold for μ_t since $\|\nu_t - \nu_0\|_2 = \|\psi \circ \mu_t - \psi \circ \mu_0\|_2$.

C Degeneracy Speed Analysis

We analyze the system degeneracy $\|\mu_t - \mu_0\|_2$ (when the candidate pool \mathcal{M} is finite) and $\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)|$ (when \mathcal{M} is infinite) asymptotically in the order of t . As $t \rightarrow \infty$, $\mu_t \rightarrow \infty$ for the selected items. Thus the selected items are asymptotically almost surely clicked, and hence for any such item a , $|\mu_t(a) - \mu_0(a)| \approx \delta(a) \cdot T_a(t)$.

Finite Candidate Pool.

$$\|\mu_t - \mu_0\|_2 \approx \sqrt{\sum_a [\delta(a) \cdot T_a(t)]^2}. \quad (17)$$

In the *Random* recommender system, $T_a(t) \approx tl/m$ and thus

$$\|\mu_t - \mu_0\|_2 \approx \sqrt{\sum_a \delta(a)^2 (l/m)} \cdot t. \quad (18)$$

Both the *Oracle* and the *Optimal Oracle* have a fixed set of items \mathcal{S} that they keep selecting. Thus $T_a(t) \approx t$ for any item $a \in \mathcal{S}$ and

$$\|\mu_t - \mu_0\|_2 \approx \sqrt{\sum_{a \in \mathcal{S}} \delta(a)^2} \cdot t. \quad (19)$$

For both the *UCB* and *TS* we have

$$\|\mu_t - \mu_0\|_2 \approx \left(\sqrt{\sum_a \delta(a)^2 / \sqrt{m^*}} \right) \cdot t, \quad (20)$$

where m^* are the number of close to optimal arms. Therefore all five models have the degeneracy quantity $\|\mu_t - \mu_0\|_2$ converge to a linear function of t .

Infinite Candidate Pool. Since we sample $\delta(a)$ from a bounded interval $[-b, b]$, asymptotically $\sup_{a \in \mathcal{M}} \delta(a) \approx b$

$$\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)| \approx \sup_{a \in \mathcal{M}} \delta(a) \cdot T_a(t) \approx bT_{a^*}(t)$$

where a^* is an item with $\delta(a^*) \approx b$.

In the *Random* recommender system, $T_a(t) \approx \text{const}$ and thus

$$\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)| \approx \text{const}. \quad (21)$$

The *Oracle* has a fixed set of items \mathcal{S} that it keeps selecting. Thus $T_a(t) \approx t$ for any item $a \in \mathcal{S}$ and

$$\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)| \approx \max_{a \in \mathcal{S}} \delta(a) \cdot t. \quad (22)$$

The *Optimal Oracle* model will pick all items with $\delta(a) \approx b$ and $T_a(t) \approx ct$ for some constant c .

$$\sup_{a \in \mathcal{M}} |\mu_t(a) - \mu_0(a)| \approx bc \cdot t. \quad (23)$$

For UCB and Thompson sampling the situation is less clear. When growth of the candidate pool is linear, then UCB will spend most of its time exploring new items and consequently degeneration is very slow. Thompson sampling will continue to play degenerate items with reasonable probability and so degeneration speed will be larger than for UCB. Precisely quantifying the rates of degeneration depends in a complicated way on the details of the model.

D Graphs

Figure 7 shows the echo chamber and filter bubble effects of three baseline models *Random*, *Oracle* and *Optimal Oracle*. Except for the *Random Model*, the other two models quickly start to only serve a selective few items in the candidate pool. Although all models cause user interest degeneration (see also Figure 2), the degeneration speeds are quite different (*Optimal Oracle* > *Oracle*, *TS*, *UCB* > *Random Model*). The *Oracle*, *TS* and *UCB* optimize based on μ_t and so we see a positive degenerative dynamics for μ_t . The *Optimal Oracle* optimizes on the degeneration speed δ directly and not on μ_t so we see both a positive and negative degeneration in μ_t . The *Random Model* also drifts μ_t in both directions, but at a much slower rate.

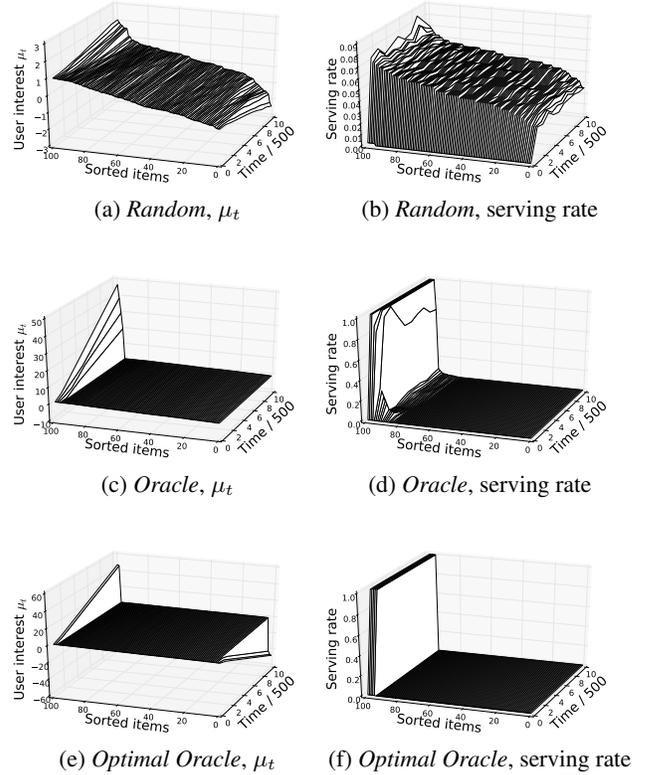


Figure 7: Echo chamber and filter bubble effect of *Random*, *Oracle* and *Optimal Oracle* models. Plotted are the user interest μ_t and serving rate of each of the 100 items in the candidate pool over 5000 steps. The x-axis marks every 500 time steps. The y-axis is the item index from 1 to 100. The z-axis is the corresponding user interest μ_t or the serving rate.