
Human Interaction with Recommendation Systems

Sven Schmit
Stanford University

Carlos Riquelme
Google Brain

Abstract

Many recommendation algorithms rely on user data to generate recommendations. However, these recommendations also affect the data obtained from future users. This work aims to understand the effects of this dynamic interaction. We propose a simple model where users with heterogeneous preferences arrive over time. Based on this model, we prove that naive estimators, i.e. those which ignore this feedback loop, are not consistent. We show that consistent estimators are efficient in the presence of myopic agents. Our results are validated using extensive simulations.

1 INTRODUCTION

We find ourselves surrounded by recommendations that help us make better decisions. However, relatively little work has been devoted to the understanding of the dynamics of such systems caused by the interaction with users. This work aims to understand the dynamics that arise when users combine the recommendations with their own preference when making a decision.

For example, a user of Netflix uses their recommendations to decide what movie to watch. However, this user also has her own beliefs about movies, e.g. based on artwork, synopsis, actors, recommendations by friends, etc. The user thus combines the suggestions from Netflix with her own preferences to decide what movie to watch. Netflix captures data on the outcome to improve its recommendations to future users. Of course, this pattern is not unique to Netflix, but observed more broadly; across all platforms that use recommendations.

Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS) 2018, Lanzarote, Spain. PMLR: Volume 84. Copyright 2018 by the author(s).

A first requirement for any estimator is consistency; however it is not clear that in the presence of human interaction naive estimators are consistent. Indeed, we show that simple estimators can easily be fooled by the selection effect of the users. We propose to measure performance by adapting the notion of regret from the multi-armed bandit literature. Using this metric, we show that naive estimators suffer linear regret; even with ‘infinite data’ the performance per time step is bounded away from the optimum.

Using the notion of regret is useful as it also allows us to quantify the efficiency of estimators. New users and items with little to no data constantly arrive and thus a recommendation system is always in a state of learning. It is therefore important that the system learn efficiently from data. While this might sound like the well-known *cold-start* problem, that is not the focus of this work; Rather than providing recommendation solutions for users in the absence of data, we focus on quantifying how quickly an algorithm obtains enough data to make good recommendations. This is more akin to the social learning and incentivizing exploration literature than work on the cold-start problem.

1.1 Main results

From a technical standpoint, this paper provides a dynamical model that captures the dynamics of users with heterogeneous preferences, while abstracting away the specifics of recommendation algorithms. In the first part of this work, we show that there is a severe selection bias problem that leads to linear regret. Second, we show that when the algorithm uses unbiased estimates for items, ‘free’ exploration occurs and we recover the familiar logarithmic regret bound. This is important because inducing agents to explore is difficult from both a statistical and strategic point of view. We validate our claims using simulations with feature-based and low-rank methods.

It is important to note that the focus of this work is to provide a simplified framework that allows us to reason about the dynamic aspects of recommendation systems. We do not claim that the model nor the as-

sumptions are a perfect reflection of reality. Instead, we believe that the model we propose provides an excellent lens to better understand vital aspects of recommendation systems.

1.2 Related work

This work roughly intersects with three separate fields of study. *Recommendation systems* [Adomavicius and Tuzhilin, 2005] have attracted much attention. In particular, much research has focused on new methods that treat the data as fixed, rather than dynamic. There has been less work on selection bias, which was first demonstrated by Marlin [2003], and subsequent work [Amatriain et al., 2009, Marlin et al., 2007, Steck, 2010]. Rather than modeling user behavior directly, they impose the statistical assumption of a covariance shift; the distribution of observed ratings is not altered by conditioning on the selection event, but five star ratings are more likely to be observed. More recently, Schnabel et al. [2016] and Joachims et al. [2017] link the bias from covariance shifts to recent advances in causal inference. Mackey et al. [2010] combine matrix factorization and latent factor models to capture heterogeneity in interactions and context.

The different approach of this work is reminiscent of the work on *social learning* [Chamley, 2004, Smith and Sørensen, 2000], where agents learn about the state of the world by combining their private signals with observations of actions (but not necessarily outcomes) of others. The work of Ifrach et al. [2014] is closest related to our setup. They discuss how consumer reviews converge on the quality of a product, given diversity of preferences under a reasonable price assumption. However, the work in social learning focuses on users interacting with a single item. This seemingly minor minor difference leads to completely different dynamics.

Finally, we can relate our work on exploration to the multi-armed bandit literature [Bubeck and Cesa-Bianchi, 2012]. In particular, there has been prior work on human interaction with multi-armed bandit algorithms: for example, how a system can optimally induce myopic agents to explore [Kremer et al., 2013] by using payments [Frazier et al., 2014] or by the way the system disseminates information [Papanastasiou et al., 2014, Mansour et al., 2015, 2016]. Similar to those works, we use the regret framework to analyse a system with interacting agents. Because in our model agents have heterogeneous preferences, we show that agents do not need to be incentivized to explore. Recent work by Bastani et al. [2017], Qiang and Bayati [2016] consider natural exploration in contextual bandit problems, and show that a modified greedy algorithm performs well. While their motivation is

different, the results are similar to ours. There has also been work on ‘free exploration’ in auction environments [Hummel and McAfee, 2014].

1.3 Organization

In the next section, we introduce our model. In Sections 3 and 4 we focus on the issues of consistency and efficiency, respectively. We illustrate our results with simulations in Section 5 before concluding.

2 MODELING HUMAN-ALGORITHM INTERACTION

In this section, we propose a model for the interaction between the recommendation system (*platform*) and users (*agents*). Each user selects one of the items the platform recommends, and reports their experience to the platform by providing a rating as feedback. The platform uses this feedback to update the recommendations for the next user.

More formally, we assume there are K items, labeled $i = 1, \dots, K$, and each item has a distinct, but unknown, quality $Q_i \in \mathbb{R}$. This aspect models the vertical differentiation between items and it is the task of the platform to estimate these qualities. For notational convenience, we assume $Q_1 > Q_2 > \dots > Q_K$. At every time step $t = 1, \dots, T$ a new user arrives and selects one of the K items. To do so, the user receives a private preference signal $\theta_{it} \sim F_i$ for each item, drawn from a preference distribution F_i which we make precise later. The value of item i for user t is

$$V_{it} = Q_i + \theta_{it} + \varepsilon_{it} \quad (1)$$

where ε_{it} is additional noise drawn independently from a noise distribution E with mean 0 and finite variance $\sigma_E^2 < \infty$. To aid the agents, the platform provides a recommendation score s_{it} , aggregating the feedback from previous agents. The agent uses her own preferences, along with the score, to select item a_t according to

$$a_t = \arg \max_i s_{it} + \theta_{it}. \quad (2)$$

Hence, we make the assumption that the agent is boundedly rational and uses s_{it} as a surrogate for the quality. Abusing notation, we write

$$V_t = V_{a_t t} = Q_{a_t} + \theta_{a_t t} + \varepsilon_{a_t t} \quad (3)$$

for the value of the chosen item for agent t . After the agent selects item a_t and observes the value V_t , the platform queries for feedback W_t from the user. For example the platform can ask the user to provide the value of the item as a rating, in which case $W_t =$

V_t . Note that the private preferences θ_t of the agent remain hidden. The platform uses this feedback to give recommendations to future users. In particular, we require s_{it} to be measurable with respect to past feedback, that is $\sigma\{a_\tau, W_\tau : \tau < t\}$.

We measure the performance of a recommendation system in terms of (pseudo-)regret:

$$R_T = \sum_{t=1}^T \max_i (Q_i + \theta_{it}) - (Q_{a_t} + \theta_{a_t t}), \quad (4)$$

which sums the difference between the expected value of the best item and the expected value of the selected item.¹ We note that if scores $s_{it} \equiv Q_i$ for all t , the regret of such platform would be 0, as each user selects her optimal action using equation (2).

2.1 Preferences, values and personalization

We use this section to expand on the motivation of the proposed model. The value for item i at time t consists of three parts (see equation (1)). First, the intrinsic quality Q_i can be seen as the mean quality across users. In our theoretical analysis we treat this as a constant to be estimated such that we are able to disentangle the model fitting from the dynamics of interaction. In this simple setting, one should view it as a vertical differentiator between items. Taking hotels as example, it could model quality of service and cleanliness, where a common ranking across agents is sensible. The intrinsic qualities Q_i can be replaced by more complicated models, for example based on feature based regression methods, or matrix factorization methods. Indeed, in Section 5.2.3 we provide simulation results where we replace Q_i with a low rank matrix factorization model.

The second term in the value equation, θ_{it} , models horizontal differentiation across agents. In our simplified model agents only arrive once, and thus this also covers different contexts. For example, one traveler prefers a hotel on the waterfront, while another prefers a hotel downtown, and yet a third prefers staying close to the convention center. While these hypothetical hotels could have the same quality, the value for users differs, in ways known to the user. However, the intrinsic quality of these properties are unknown to these users.

All in all, the value of item i for agent t is drawn from a distribution with mean Q_i , and where the variance consists of a part that is known to the user (θ_{it}) and a part that is unknown to both platform and user (ε_{it}).

¹ Unlike the traditional bandit setting, there is no single best item. Rather, different users might have different optimal items.

In section 5 we investigate how well the theoretical results carry over to more general models.

One could argue that personalization methods (i.e. replacing Q with more sophisticated models) supersede the need for idiosyncratic preferences θ_{it} , as these preferences can be captured by those models. However, we argue that in most cases this factor cannot be eliminated. Every recommendation system is constrained in terms of the *quantity* and *quality* of the data it is based on. First, a user only interacts with a system so often, and that limits the amount of personalization that models can achieve. Second, recommendation systems often have access to only weak features, and some aspects of user preferences and contexts, such as taste or style, can be difficult to capture. Together, these constraints make it difficult to fully model users preferences, hence the need to explicitly model the unobserved preferences to get a deeper understanding of the dynamics of recommendations systems.

2.2 Incentives

We note that the agents in our model are boundedly rational: their behavior is not optimal, and in particular ignores the design of the platform. Experimentally, there has been abundant evidence of human behavior that is not rational [Camerer, 1998, Kahneman, 2003]. Simple heuristics of user behavior have been used by others in the social learning community. Examples include learning about technologies from word-of-mouth interactions [Ellison and Fudenberg, 1993, 1995] and modeling persuasion in social networks [Demarzo et al., 2003]. The combination of machine learning and mechanism design with boundedly rational agents is explored by Liu et al. [2015].

The behavior of the user in our model implicitly relies on three assumptions:

1. The user is naive; she believes the scores supplied by the platform are unbiased estimates of the true quality.
2. The user is myopic; she selects the item that seems best for her.
3. The user has incentives to give honest feedback.

The first assumption seems unrealistic if the platform abuses this power to dictate exploration, which does not align with the myopic behavior. However, in Section 4 we show that there is no need for such aggressive exploration from the platform to obtain order-optimal performance. We also note that if the platform outputs the true qualities Q_i , then the selection rule (2) is optimal for a myopic agent. Finally, it is not obvious why a myopic user would leave feedback. While

we do not explicitly model returning users, we argue that in general a user is motivated to leave feedback because it leads to better recommendations for her in the future.

3 CONSISTENCY

In this section, we analyse the performance of standard algorithms, that is, scoring processes that do not take into account that agents have private preferences, and base the scores on empirical averages. This does include algorithms that trade-off exploration and exploitation, such as variants of UCB [Auer et al., 2002] and Thompson Sampling [Russo and Roy, 2016]. We focus on the Bernoulli preferences model, though in Section 5 we empirically demonstrate that different preference distributions lead to similar outcomes.

First we define the set of agents before time t that have selected item i by

$$S_{it} = \{\tau < t : a_\tau = i\}. \quad (5)$$

We also define \bar{V}_{it} to denote the empirical average of item i up to time t :

$$\bar{V}_{it} = \frac{1}{|S_{it}|} \sum_{\tau \in S_{it}} V_\tau \quad (6)$$

where we use $\bar{V}_{it} = 0$ when $S_{it} = \emptyset$. We want to show that the system suffers linear regret when the platform uses any scoring mechanism for which scores converge to the empirical average of the observed values. This means that the system never converges to an optimal policy; rather a constant fraction of users are misled into perpetuity. To make this rigorous, we define the notion of *mean-converging* scoring process.

Definition 1. *A scoring process that outputs scores s_{it} for item i at time t is mean-converging if*

1. s_{it} is a function of $\{V_\tau : \tau \in S_{it}\}$ and t .
2. $s_{it} \rightarrow \bar{V}_{it}$ almost surely if $\liminf_t \frac{|S_{it}|}{t} > 0$ almost surely.

In words, the score only depends on the observed outcomes for this particular item, and if we observe a linear number of selections of arm i , then the score converges to the mean outcome. Trivially, this includes using the average itself as score, $s_{it} = \bar{V}_{it}$, but this definition also includes well known methods that carefully balance exploitation with exploration, such as versions of UCB and Thompson Sampling.

From the previous section, we know that, ideally, the scores supplied to the user converge to the quality of

the item, $s_{it} \rightarrow Q_i$, as more users select item i . We say that the scores are *biased* if this is not the case:

$$s_{it} \not\rightarrow Q_i \quad \text{as} \quad |S_{it}| \rightarrow \infty. \quad (7)$$

The next proposition shows that mean-converging scoring processes lead to linear regret, because these scores are generally biased. We only show this result for when preferences are drawn from Bernoulli distributions as this simplifies the analysis significantly. In Section 5 simulations show that linear regret is observed under a variety of preference distributions. Under the Bernoulli model, it is needed that the gap between qualities is ‘small’, though we show that this condition is rather weak.

Proposition 1. *When $\theta_{it} \sim \text{Bernoulli}(p)$ for all i, t , if*

$$\Delta = Q_1 - Q_2 < \frac{(1-p)^K}{(1-p)^K + p} \quad (8)$$

and s_{it} is mean-converging, then

$$\limsup_{t \rightarrow \infty} \frac{R_t}{t} \geq c \quad (9)$$

for some $c > 0$.

The proof of this proposition can be found in the supplemental material. The intuition behind the result is that the best ranked item is selected by users that do not necessarily like it that much, while other items are only selected by users who really love it. Therefore, the ratings of the best ranked item suffers relative to others.

We note that for $K = 100$ and $p = 1/50$, the condition requires $\Delta < 0.86$. More generally, in the relevant regime where $p < \frac{\log(K)}{2K}$, the condition is satisfied if $\Delta < 0.7$ for all K . We also note that the linear regret we obtain is not caused by the usual exploration/exploitation trade-off, but rather the estimators being biased.

There is no bias result for general preference distributions as it is possible to cherry pick distributions in such a way that biases cancel each other out exactly. Furthermore, the magnitude of the problem depends crucially on the variance in user preference relative to the differences in qualities.² However, in Section 5 we provide simulations with a variety of preference distributions that suggest that bias is not an artifact of our assumptions.

3.1 Unbiased estimates

Naturally, a first attempt to improve the linear regret is aimed at obtaining unbiased versions of the naive averaging. We now sketch a few such approaches.

² In the limiting scenario of no variance in preferences, we already know that there is no bias either.

3.1.1 Randomization

Researchers can avoid selection bias in experimental studies by randomizing treatments, and we can employ the same approach here. Instead of the user choosing an action, the platform assigns matches between users and items. Note that pure randomization is not needed, user ratings are unbiased as long as the selection of items is independent from private preferences.

Just like randomized control studies, this approach is often infeasible or prohibitively expensive, e.g. a platform cannot force the user to select a certain hotel or restaurant. However, small scale user experimentation can potentially inform the platform on the magnitude and effects of idiosyncratic preferences and the bias it induces.

3.1.2 Algorithmic approach

Another option is to consider algorithmic approaches to obtain consistent estimators. In the case of Bernoulli preferences, it is possible to obtain unbiased estimates from data. However, this approach does not generalize to other preference distributions, let alone to the case where we do not know the underlying preference distributions.

3.1.3 Changing the feedback model

Given the difficulty of debiasing feedback algorithmically, we briefly discuss a third alternative. The traditional type of question ‘How would you rate this item?’ asks for an absolute measure of satisfaction, which corresponds to directly probing for V_t in our model. If we ask how the chosen item compares to the expectation, we ask for a relative measure of feedback, approximating $V_t - (s_{a_t} + \theta_{a_t})$. An example of such prompt could be ‘How does this item compare to your expectation?’ This way, we can uncover an unbiased estimate of Q_{a_t} . Importantly, it does not require any distributional assumptions on the form of the preferences. Whether such relative feedback works in practice would require a thorough empirical study, which is beyond the scope of this work. We also note that this approach does not work when platforms collect implicit feedback.

4 EFFICIENCY

From the previous section we know that naive scoring mechanisms are inconsistent and lead to linear regret. We now focus on the efficiency of consistent estimators, and we assume we have access to unbiased feedback from now on. But, this is not necessarily sufficient to guarantee good performance. The multi-armed bandit literature suggests that algorithms with small regret require a careful balance between exploration and

exploitation.

In particular, that means that the system needs to obtain data on every item in order to provide useful scores to the users. However, myopic agents have no interest in assisting the platform with exploration. In this section we address the problem of exploration in the proposed model. As opposed to the research mentioned in the introduction, we deal with agents with heterogeneous preferences. It seems natural that these heterogeneous agents help the system explore, but it is not obvious to what extent this helps. We show that because of this diversity in preferences, the free exploration leads to optimal performance (up to constants); we recover the standard logarithmic regret bound from the bandit literature. This means that there is little need for a platform to implement a complicated exploration strategy, and incentives naturally align much better than in the settings of previous work.

4.1 Formal result

We assume the qualities Q_i are bounded, and without loss of generality we can assume they are bounded in $[0, 1]$, and we assume access to unbiased feedback from the user. That is, the feedback at time t for chosen item i_t is $\tilde{V}_t = Q_{i_t} + \varepsilon_t$. However, we no longer require that private preference distributions are Bernoulli. Let \bar{V}_{it} denote the average of the values observed of item i by time t

$$\bar{V}_{it} = \frac{1}{|S_{it}|} \sum_{\tau \in S_{it}} \tilde{V}_\tau \quad (10)$$

where $S_{it} = \{\tau < t : a_\tau = i\}$. We now consider the scoring algorithm that clips the value onto $[0, 1]$: $s_{it} = \max(0, \min(1, V_{it}))$.

Then, if the private preferences have sufficiently large variance, made precise in the theorem statement, then exploration is guaranteed and the platform suffers logarithmic regret. Let Δ_{\min} be the smallest gap in qualities $\Delta_{\min} = \min_{i,j} |Q_i - Q_j|$.

The following result shows that empirical averaging is enough to get an order optimal (pseudo)-regret bound with respect to the total number of agents T .

Proposition 2. *If for all i , F_i are such that $\mathbb{P}(\theta_{it} > 1) > \gamma'$ and $\mathbb{P}(\theta_{it} \leq 0) \geq \gamma$ Then*

$$\mathbb{E}[\text{regret}(T)] \leq \left(\frac{16\sigma^2}{\Delta_{\min}} + \Delta_{\min} \right) K + \frac{32\alpha\sigma^2 K (\log(T) - \log(\Delta_{\min}) + \log(2))}{\Delta_{\min}^2 C} \quad (11)$$

where $C = \gamma' \gamma^{K-1}$.

The proof can be found in the supplement. The main idea is that we can show every arm is chosen suffi-

ciently often initially, and then concentration inequalities ensure good performance after an initial learning phase.

Corollary 3. *Suppose F_i is a symmetric distribution around 0 such that $\mathbb{P}(\theta_{it} > 1) > \gamma$, then the theorem applies with $C \geq \gamma 2^{1-K}$.*

Also note that the theorem applies for Bernoulli preferences, where $C = p(1-p)^{K-1}$. The small value of C leads to a large leading constant. Simulations suggest that performance is much better in practice.

4.2 Remarks

The main take-away from this result is not so much the specific bound, but rather the practical insight that it, and its proof, yield. The intuition is that initially estimates of quality are poor. Therefore, it takes some time and luck for users with idiosyncratic preferences to try these items. As estimates improve, however, most agents are drawn to their optimal choice. Since these choices differ across agents, the platform gets to learn efficiently about all items without incurring a regret penalty.

The practical consequence of this observation is that to improve the performance, the designer of a recommendation system should focus on simple ways to make new items, or more generally items with few observations, more likely to be chosen. We can achieve this by highlighting new arrivals. A good example is Netflix, which clearly displays a ‘Recently Added’ selection.

5 SIMULATIONS

In this section we empirically demonstrate that the theoretical results derived in the previous sections hold much more broadly. First, we focus on verifying the results from our simplified model. Thereafter, we consider more advanced personalization models using feature-based and low rank methods, where we investigate the dynamics with private preferences.³

5.1 Simulations of regret

Before considering more advanced methods, we simulate our model using different preference distributions and plot the cumulative regret over time. We run 50 simulations with 5000 time steps and $K = 50$ items across four preference distributions with randomly drawn parameters. We then compare biased and unbiased algorithms based on empirical averages. Figure 1 shows the cumulative regret paths for each

³ The code to replicate the simulations is publicly available at https://github.com/schmit/human_interaction.

of these simulations. The qualities were drawn from the uniform distribution over $[0, 1]$. For the preference distributions F_i for item i , we used

- Bernoulli distribution with $p_i \sim U[0, \frac{2 \log(K)}{3K}]$.
- Normal distribution with $\mu_i = 0$ and $\sigma_i \sim U[0, 1]$.
- Exponential distribution with scale $\lambda_i^{-1} \sim U[0, 1]$.
- Pareto distribution with shape $\alpha_i \sim U[2, 4]$.

These are chosen such that the variance in preference and qualities is roughly similar. A clear pattern emerges; In all cases, the (biased) empirical averages lead to linear regret, not just for the Bernoulli model covered by Proposition 1. Second, we note the unbiased scores lead to much better results regardless of the preference distribution, in line with Proposition 2.

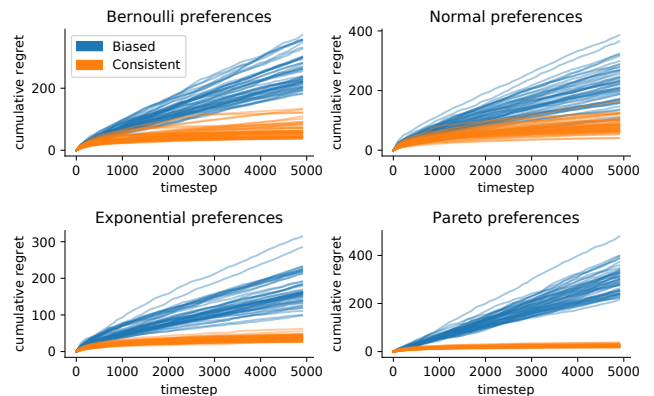


Figure 1: These plots show the cumulative regret plotted against time for both the naive empirical averages in blue, and the unbiased averages in orange.

5.2 Personalization methods

We now focus on methods that provide more personalized recommendations. Because estimating such models is more complicated and computationally intensive, we simplify the dynamics of our simulations to a two-staged approach. We then use this approach to experiment with a feature-based and a low-rank approximation approach to personalization based on synthetically generated data.

5.2.1 Two-staged simulations

In our original setup, the platform updates its scoring rule after every observation. This is impractical when dealing with more sophisticated models. Instead we first collect a set of observations using a fixed scoring algorithm (a *training set*), and fit the scoring algorithm once to this training data. We then use this

trained algorithm to generate a new set of observations (a *test set*), again without updating the algorithm in between observations. The test set is used to measure the performance of the fitted algorithm. Instead of using regret, we measure performance by directly computing the average rating on the second dataset generated by our trained algorithm. Note that it is possible to iterate generating data and fitting models multiple times before generating a test set.

5.2.2 Ridge regression

In this section we discuss a feature-based model of personalization where the rating is assumed to be a linear function of observed covariates.

The model In the feature-based setting, each item has an unknown parameter vector w_i and each user-item pair has an observed feature vector x_{it} . The value of item i for user t then becomes

$$V_{it} = Q_i + x_{it}^T w_i + \theta_{it} + \varepsilon_{it}. \quad (12)$$

Furthermore, we also parametrize θ_{it} in terms of x_{it} , such that

$$\theta_{it} \sim \mathcal{N}(x_{it}^T \tilde{w}, \sigma_\theta) \quad (13)$$

where \tilde{w} is another unknown parameter vector. After generating the training set using a fixed scoring rule, we use ridge regression to regress the reported ratings for each item, which leads to estimates \hat{Q}_i and \hat{w}_i for each item. These are then used as scoring rule: $s_{it} = \hat{Q}_i + x_{it}^T \hat{w}_i$.

Simulation details There are $n = 100$ items, and $w_i \in \mathbb{R}^p$ where $p = 20$. We generate $Q_i \sim \mathcal{N}(0, 1)$ and $w_{ij}, \tilde{w}_{ij} \sim \mathcal{N}(0, 1/\sqrt{p})$ independently. The elements of the feature vectors are generated independently following $x_{ijt} \sim \mathcal{N}(0, 1)$. The error term is drawn according to $\varepsilon_{it} \sim \mathcal{N}(0, 1)$, and we set $\sigma_\theta = 0.1$. We generate $10np = 20000$ observations.

The training sets are generated using four different scoring rules:

1. Using the *oracle* scoring rule: $s_{it} = Q_i + x_{it}^T w_i$, which leads to perfect recommendations.
2. Using the oracle scoring rule and unbiased ratings $V_{it} - \theta_{it}$.
3. Using randomly selected items, hence the user has no choice.
4. We iterate steps one and two twice, where we first use randomly selected items, then fit a Ridge regression to estimate the parameters, and use these to generate the test set: $s_{it} = \hat{Q}_i + \hat{x}_{it}^T \hat{w}_i$.

This last training set allows us to better understand how the system evolves over time.




Figure 2: We compare performances of recommendation systems based on different training data in a feature-based model on the left and matrix factorization on the right. Green corresponds to random selections in the training set, orange to oracle selections with unbiased feedback, red to the iterated model initially based on random selections, and blue is based on biased oracle ratings.

Results The average values of the selected items in the test set are plotted in the left plot of Figure 2. The two dotted lines provide useful benchmarks: the top line shows the performance of the oracle scoring rule, which upper bounds the performance. The bottom line shows the performance in the absence of a scoring rule, that is $s_{it} = 0$ for all i, t . Finally, note that random selections lead to an average rating of 0.

We note that the best performing recommendations are given by the model trained on random data (green); these are close to the performance of the oracle. Unbiased ratings based on the oracle (orange) perform a bit worse due to a feedback loop. The model is only trained on ‘good’ selections and this leads to a degradation of performance. We also see that the iterated model (red) that was initially trained on random selections performs worse than the randomly generated data, suggesting that the quality deteriorates over time. Finally, the model trained on oracle data (blue) performs much worse than all the models, and does not perform much better than the ‘no-score algorithm’ that does not provide recommendations.

5.2.3 Matrix factorization

In this section we investigate the dynamics of private preferences that are low rank. We use the same two-stage approach as before, where we first use a fixed scoring rule to generate a training set, fit our model, and use the fitted scoring rule to generate a test set to measure performance.

The model The low-rank models assumes that the

value for item i by user j follows the model

$$v_{ij} = a_i + b_j + u_i^T v_j + x_i^T y_j + \varepsilon_{it} \quad (14)$$

where u_i and v_j are (hidden) q -dimensional vectors modeling interaction between user and item, and a_i and b_j are terms modeling overall differences between users and items. Similarly, x_i and y_j are q -dimensional vectors that combine the private preference $\theta_{it} = x_i^T y_j$. Note that, unlike in previous settings, here we observe the same user multiple times. The recommendation system provides score s_{ij} for user i and item j and the user selects item

$$\arg \max_j s_{ij} + x_i^T y_j \quad (15)$$

and reports her value V_{ij} . We use alternating least squares [Koren et al., 2009] to estimate a , b , u and v .⁴

Simulation details As in the feature-based simulation, we generate four training sets, one based on an oracle, one based on an oracle with unbiased ratings, one based on random selections, and finally an iterated version of the random selections process, where we fit a model to the random selections data and use that model to generate training data.

We simulate 2000 users and 500 items, with rank $q = 4$. Entries of u , v , x and y are independent Gaussians with variance $1/q$. To reduce variance, the error term has a small variance, $\varepsilon_{ij} \sim N(0, 0.01)$ and for both training and test sets each user rates 40 items. We run alternating least squares with rank $2q$ and varying regularization.

Results The two dotted lines denote the same benchmarks as before. Again, we notice that recommendations trained on random data perform best, but this time the difference is much more pronounced. The recommendations based on perfect recommendations (blue and orange) perform a lot worse. In fact, they do barely better than not recommending items at all and having users base their choice solely on their own preference signals. Part of the degradation in performance seems to be caused by a feedback loop; the observations are not randomly sampled. We also notice a much stronger degradation in performance of the iterated model (red). This suggests that the dynamic nature of recommendation systems affect matrix factorization methods more severely than the simpler linear model from the previous section.

6 DISCUSSION

In this work, we introduce a model for analyzing feedback in recommendation systems. We propose a sim-

⁴ We ensure that users do not rate the same item in both the training and test set.

ple model that explicitly looks at heterogeneous preferences among users of a recommendation system, and takes the dynamics of learning into account. We then consider the consistency and efficiency of natural estimators in this model. Recent work has focused on exploration, or efficiency, with selfish agents. On the one hand, preferences lead to inconsistent estimators if this aspect is not taken into account. On the other hand, we also show that there is an upside to heterogeneous preferences; they automatically lead to efficiency. Using simulations, we demonstrate that these phenomena persist when we use more sophisticated recommendation methods, such as matrix factorization.

6.1 Future work

There are several directions of further research. Our simplified model does not capture all aspects of recommendation systems. The most interesting aspect is that, in practice, users only observe a limited set of recommendations, rather than the entire inventory. This can lead to an inefficiency in the rate of exploration, and requires further study.

Our model and simulations show that consistency of models is an issue that is difficult to resolve. We believe that progress can be made. Theoretically, one possible avenue is to also model the selection process directly and combine it with the model for outcomes. Empirically, by large scale studies that test the effects of human interaction on estimators.

6.2 The bigger picture

We believe that this work has raised fundamental and important issues relating the interaction between machine learning systems and the users interacting with them. Algorithms not only consume data, but in their interaction with users also create data, a much more opaque process but equally vital in designing systems that achieve the goals we set out to achieve. There is still a lot of room for improvement by gaining a better understanding of these dynamics.

7 Acknowledgements

The authors would like to thank Ramesh Johari, Vijay Kamble, Brad Klingenberg, Yonatan Gur, Peter Lofgren, Andrea Locatelli, and anonymous reviewers for their suggestions and feedback. This work was supported by the Stanford TomKat Center, and by the National Science Foundation under Grant No. CNS-1544548. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.*, 17:734–749, 2005.
- X. Amatriain, J. M. Pujol, and N. Oliver. I like it... i like it not: Evaluating user ratings noise in recommender systems. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 247–258. Springer, 2009.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- H. Bastani, M. Bayati, and K. Khosravi. Exploiting the natural exploration in contextual bandits. *arXiv preprint arXiv:1704.09011*, 2017.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *CoRR*, abs/1204.5721, 2012.
- C. Camerer. Bounded rationality in individual decision making. *Experimental economics*, 1(2):163–183, 1998.
- C. Chamley. *Rational Herds: Economic Models of Social Learning*. Rational Herds: Economic Models of Social Learning. Cambridge University Press, 2004. ISBN 9780521530927. URL <https://books.google.com/books?id=2dgb0h6VE9YC>.
- P. M. Demarzo, D. Vayanos, J. Zwiebel, N. Barberis, G. Becker, J. Bendor, L. Blume, S. Board, E. Dekel, S. Dellavigna, D. Duffie, D. Easley, G. Ellison, S. Gervais, E. Glaeser, K. Judd, D. Kreps, E. Lazear, G. Loewenstein, L. Nelson, A. Neuberger, M. Rabin, J. Scheinkman, A. Schoar, P. Sorenson, P. Veronesi, and R. Zeckhauser. Persuasion bias, social influence, and unidimensional opinions. 2003.
- G. Ellison and D. Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101(4):612–643, 1993.
- G. Ellison and D. Fudenberg. Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, 110(1):93–125, 1995.
- P. I. Frazier, D. Kempe, J. M. Kleinberg, and R. Kleinberg. Incentivizing exploration. In *SIGECOM*, 2014.
- P. Hummel and R. P. McAfee. Machine learning in an auction environment. In *WWW*, 2014.
- B. Ifrach, C. Maglaras, and M. Scarsini. Bayesian social learning with consumer reviews. *SIGMETRICS Performance Evaluation Review*, 41:28, 2014.
- T. Joachims, A. Swaminathan, and T. Schnabel. Unbiased learning-to-rank with biased feedback. *CoRR*, abs/1608.04468, 2017.
- D. Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *The American economic review*, 93(5):1449–1475, 2003.
- Y. Koren, R. M. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42, 2009.
- I. Kremer, Y. Mansour, and M. Perry. Implementing the "wisdom of the crowd". In *SIGECOM*, 2013.
- T.-Y. Liu, W. Chen, and T. Qin. Mechanism learning with mechanism induced data. In *AAAI*, pages 4037–4041, 2015.
- L. W. Mackey, D. J. Weiss, and M. I. Jordan. Mixed membership matrix factorization. In *ICML*, 2010.
- Y. Mansour, A. Slivkins, and V. Syrgkanis. Bayesian incentive-compatible bandit exploration. *CoRR*, abs/1502.04147, 2015.
- Y. Mansour, A. Slivkins, V. Syrgkanis, and Z. S. Wu. Bayesian exploration: Incentivizing exploration in bayesian games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 661–661, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0. doi: 10.1145/2940716.2940755. URL <http://doi.acm.org/10.1145/2940716.2940755>.
- B. M. Marlin. Modeling user rating profiles for collaborative filtering. In *NIPS*, 2003.
- B. M. Marlin, R. S. Zemel, S. T. Roweis, and M. Slaney. Collaborative filtering and the missing at random assumption. In *UAI*, 2007.
- Y. Papanastasiou, K. Bimpikis, and N. Savva. Crowdsourcing exploration. *History*, 2014.
- S. Qiang and M. Bayati. Dynamic pricing with demand covariates. 2016.
- D. Russo and B. V. Roy. An information-theoretic analysis of thompson sampling. *Journal of Machine Learning Research*, 17:68:1–68:30, 2016.
- T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, and T. Joachims. Recommendations as Treatments: Debiasing Learning and Evaluation. *ArXiv e-prints*, Feb. 2016.
- L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- H. Steck. Training and testing of recommender systems on data missing not at random. In *KDD*, 2010.