

Movie Recommender System for Profit Maximization

Amos Azaria¹, Avinatan Hassidim¹, Sarit Kraus^{1,2}, Adi Eshkol³, Ofer Weintraub³

Irit Netanely³

¹ Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

² Institute for Advanced Computer Studies University of Maryland, MD 20742

³ Viaccess-Orca, 22 Zarhin Street, Ra'anana 43662, Israel

{azariaa1,avinatan,sarit}@cs.biu.ac.il,

{adi.eshkol, ofer.weintraub,irit.netanely}@viaccess-orca.com

ABSTRACT

Traditional recommender systems minimize prediction error with respect to users' choices. Recent studies have shown that recommender systems have a positive effect on the provider's revenue.

In this paper we show that by providing a set of recommendations different than the one perceived best according to user acceptance rate, the recommendation system can further increase the business' utility (e.g. revenue), without any significant drop in user satisfaction. Indeed, the recommendation system designer should have in mind both the user, whose taste we need to reveal, and the business, which wants to promote specific content.

In order to study these questions, we performed a large body of experiments on Amazon Mechanical Turk. In each of the experiments, we compare a commercial state-of-the-art recommendation engine with a modified recommendation list, which takes into account the utility (or revenue) which the business obtains from each suggestion that is accepted by the user. We show that the modified recommendation list is more desirable for the business, as the end result gives the business a higher utility (or revenue). To study possible reduce in satisfaction by providing the user worse suggestions, we asked the users how they perceive the list of recommendation that they received. Our findings are that any difference in user satisfaction between the list is negligible, and not statistically significant.

We also uncover a phenomenon where movie consumers prefer watching and even paying for movies that they have already seen in the past than movies that are new to them.

1. INTRODUCTION

The main goal in designing recommender systems is usually to predict the user's wish list and to supply her with the best list of recommendations. This trend is prevalent whether we consider a social network recommending friends [2], consumer goods [14] or movies [11].

However, in most cases, the engineers that design the recommender system are hired by the business which provides the suggestions (in some cases the web-site buys a recommendation engine

from a third party - but also in these cases the site is paying for the recommendations). The business' end goal is usually to increase sales, revenues, user engagement, or some other metric. In that sense, the user is not the end customer of the recommendation system, although she sees the recommendations [13]. Still, one could argue that it is better for the business to give the user the best possible recommendations, as it will also maximize the business' profit, either in the short run (no point in giving recommendations which are not followed by the users) or at least in the long run (good recommendations make users happy).

In this paper, we provide evidence that a business may gain significantly by providing users with recommendations that may not be best from the users point of view but serve the business' needs. We provide an algorithm which uses a general recommender system as a black-box and increases the utility of the business. We perform extensive experiments with it in various cases. In particular, we consider two settings:

1. The *Hidden Agenda* setting: In this setting, the business has items that it wants to promote, in a way which is opaque to the user. For example, a movie supplier which provides movies on a monthly fee basis but has different costs for different movies, or a social network which wants to connect users who are less engaged to more engaged ones. Netflix, for instance, set a filter to avoid recommending new releases which have high costs to them [24].
2. The *Revenue Maximizing* setting: In this case the goal of the recommender system is to maximize the expected revenue, e.g. by recommending expensive items. In this setting, there is an inherent conflict between the user and the business.

To study these settings, we conducted experiments on Amazon Mechanical Turk (AMT) in which subjects were asked to choose a set of favorite movies, and then were given recommendations for another set of movies. For each recommendation, the subjects were asked if they would or would not watch the movie¹. Finally, in order to test possible reduce in satisfaction caused by tuning the recommendations to the business' needs, we asked each subject how good she feels about the recommendations she received.

This form of experimentation makes two assumptions. First, we simulate long term effects by asking users about their satisfaction from the list. It is possible that users would report that the modified list is satisfactory, but would enjoy it less subconsciously, and this would bias them against our system in the long run. Second,

¹In the Revenue Maximization setting each recommended movie also came with a price tag. We describe the experiments later in the experiments section.

we assume that asking users if they are willing to pay for a movie is the same as actually taking their money and showing them the movie. Both assumptions are common in the literature (for a comparison between hypothetical and real scenarios see [18]). We hope to integrate the algorithm in a real world system to circumvent the assumptions.

Manipulating the recommender system in order to increase revenue (or to satisfy some other hidden agenda) raises some ethical concerns. If users believe that a particular algorithm is being used (e.g. collaborative filtering), then they could be irritated if they find out that recommendations are being edited in some way. However, most businesses do not provide the specification of their recommender system (treating it as a “secret sauce”), which diminishes this concern. Furthermore, several companies (including Netflix, Walmart and Amazon) admitted human intervention in their recommender system [17], so it may well be that different companies are already tweaking their recommender systems for their own good. In this sense, an important lesson to take away from this work is “users beware”. We show that businesses garner a large gain by manipulating the system, and many companies could be tempted by this increase in revenue. In this paper we propose a method which allows businesses to mount their existing recommender system in order to increase their revenue.

An interesting phenomenon that we uncover is that subjects are more willing to pay for movies that they’ve already seen. While a similar phenomena is known for other types of consumer goods, coming across it with regards to movies is new and somewhat counter-intuitive. Interestingly, in movie watching, two opposite effects come into play: 1. Variety seeking: users want new experiences (see the survey [15] for a model of variety seeking customers). 2. Loyalty: users are loyal to brands that they know and used. It is usually assumed that in movies the former effect is far more dominant [12, 7]. Although we supply some explanation for this phenomenon, further research is required in order to fully understand it. We note that this phenomena has implications to the design of recommender systems – some systems take care not to recommend a movie which was previously watched. Instead, one may be better off recommending movies that the user saw in the past as well.

To summarize, the main contribution of this paper is twofold:

- Providing an algorithm for utility maximization of a movie supplier service, in two different settings, one with prices and the other without. This algorithm is provided along with an extensive experiment demonstrating its performance.
- Uncovering the interesting behavior of movie consumers, who prefer watching and paying for movies that they have already seen in the past (to watching and paying for movies which are new to them).

These findings have major consequences for design of recommendation systems for movie providers..

2. RELATED WORK

Models for predicting users’ ratings have been proposed that are used by recommendation systems to advise their users (See Ricci et al. [19] for a recent review). It has been shown that recommender systems, in general, are beneficial for the providing business [22]. However, most works in this realm do not explicitly try to maximize the system’s revenue, but only consider the utility of the user. The works that do try to directly increase the system’s revenue usually take a more holistic approach, changing the pricing and using the inner workings of the recommendation system. As far as we know, this is the first work which treats the recommender system as a

black box, does not change the pricing and tries to increase the revenue directly.

Pathak et al. [17] study the cross effects between sales, pricing and recommendations on Amazon books. They show that recommendation systems increase sales and cause price changes. However, the recommendation systems that they consider are price independent, and the effect on prices is indirect - items which are recommended more are bought more, which affects their price (the pricing procedure used in their data takes popularity into account). They do not consider having the price as an input to the system, and do not try to design new recommendation systems.

Chen et al. [4] develop a recommender system which tries to maximize product profitability. Chen et al. assume the usage of a collaborative filtering recommender system which, as part of its construction, provides a theoretically-based probability that a user will purchase each item. They multiply this probability by the revenue from each item and recommend the items which yield the highest expected revenue. However, in practice, many recommender systems do not rely only upon collaborative filtering (which can’t be applied to new items or when the data is sparse), but also rely on different engines (such as popularity, semantic similarity, etc.). Even a business using a pure collaborative filtering engine may not necessarily have access to (or may not want to access) the internal workings of their own recommender system. Therefore, we assume a generic recommender system which is treated as a black-box component, and dedicate most of our work to building a human model in order to predict the acceptance rate of a given item using a generic recommender system.

Das et al. [5] provide a mathematical approach for maximizing business revenue using recommender systems. However, they assume that as long as the recommendations are similar enough to the customer’s own ratings, the customer is likely to follow the recommendations. Therefore, Das et al. do not model the actual drop in user acceptance rate as the item becomes less relevant or as the item price increases, as is done in this work. Similarly, Hosanagar et al. [9] use a mathematical approach to study the conflict which a business confront when using recommender systems. On one hand the business would like to recommend items with higher revenue (margins), but on the other hand it would like to recommend items which the users are more likely to buy. Hosanagar et al. show that in order to increase its total revenue, the business must balance between these two factors. Unfortunately, neither paper provides any actual experimental evaluation with people, as is provided in this paper.

Shani et al. [23] use a discrete-state MDP model to maximize the system’s utility function, taking into account the future interactions with their users. In their work, the system may, for example, decide to recommend a game console, since, if purchased, the user is likely to purchase many games for it in the future. This approach is not applicable to our domain as we consider a single-shot recommendation. Furthermore, in the movie recommendation domain, most movies have very little impact on each other, and watching a certain movie doesn’t usually increase or decrease the probability that the user will watch a different movie. In [1] we model the long-term affect of advice given by a self-interested system on the users in route selection problems. In order for the system to maximize its long term expected revenue, we suggest that it uses what we term the “social utility” approach. However, in [1] we assume that the user must select his action among a limited number of options and the system merely recommends a certain action. Therefore the system does not act as a classic recommender system, which recommends a limited number of items from a very large corpus. Still, both of the above works may be found useful if combined with the

approach given in this paper, when considering repeated interactions scenarios.

We now turn to discuss related work to the phenomenon where movie consumers prefer watching movies that they have already seen. The marketing literature contains many examples in which an individual experiencing one type of event is more likely to experience it again. Such examples include unemployment, accidents and buying a specific product (or brand). Heckman [8] discusses two possible explanations: either the first experience changes the individual and makes the second one more likely (e.g. the individual bought a brand and liked it), or that this specific individual is more likely to have this experience (e.g. a careless driver has a higher probability of being involved in a car accident). Kamakura and Russell [10] show how to segment a market into loyal and non-loyal customers, where the loyal customers are less price sensitive and keep buying the same brand (see the works of [6] and [20] who show that even a very short purchase history data can have a huge impact, and that of [16] which shows that consuming one product from a particular company increases the probability of consuming another product from the same company).

Russel et al. [21] interviewed 23 subjects regarding repeated experiences. They describe several benefits people might have from rereading a book, re-watching a movie etc. For example, when rereading a book, the person can pay more attention to the language and the expressions used. Since people’s memory is imperfect, when re-watching a movie, people reconstruct their memory, and this reconstruction may tack back and forth between the actual memory and new knowledge acquired since. The reconstructive experience may add levels of complexity and depth.

3. PUMA

In this section we present the Profit and Utility Maximizer Algorithm (PUMA). PUMA mounts a black-boxed recommender system which supplies a ranked list of movies. This recommender system is assumed to be personalized to the users, even though this is not a requirement for PUMA.

3.1 Algorithm for Hidden Agenda Setting

In the hidden agenda setting, the movie system supplier wants to promote certain movies. Movies aren’t assigned a price. We assume that each movie is assigned a promotion value, $v(m)$, which is in $V = \{0.1, 0.2, \dots, 1\}$. The promotion value is hidden from the user. The movie system supplier wants to maximize the sum of movie promotions which are watched by the users, i.e. if a user watches a movie, m , the movie supplier gains $v(m)$; otherwise it gains nothing.

The first phase in PUMA’s construction is to collect data on the impact of the movie rank ($r(m)$) in the original recommender system on the likelihood of the users to watch a movie ($p(m)$). To this end we provide recommendations, using the original recommender system, ranked in leaps of a given k (i.e. each subject is provided with the recommendations which are ranked: $\{1, k, \dots, (n - 1) \cdot k + 1\}$ for the given subject in the original recommender system). We cluster the data according to the movie rank and, using least squared regression, we find a function that best explains the data as a function of the movie rank. We consider the following possible functions: linear, exponent, log and power (see Table 1 for function forms). We do not consider functions which allow maximum points (global or local) which aren’t at the edges, as we assume that the acceptance rate of the users should be highest for the top rank and then gradually decrease. Since these functions intend to predict the probability of the acceptance rate, they must return a value between 0 and 1, therefore a negative value returned must be set to 0 (this

Table 1: Function forms for considered functions. α and β are non-negative parameters and $r(m)$ is the movie rank.

| function | function form |
|------------------------------|--------------------------------------|
| linear (decay) | $\alpha - \beta \cdot r(m)$ |
| exponent (exponential decay) | $\alpha \cdot e^{-\beta \cdot r(m)}$ |
| log (logarithmic decay) | $\alpha - \beta \cdot \ln(r(m))$ |
| power (decay) | $\alpha \cdot r(m)^{-\beta}$ |

may only happen with the linear and log functions - however, in practice, we did not encounter this need).

Among the functions that we tested, the linear function turned out to provide the best fit to the data in the hidden agenda setting (it resulted with the greatest coefficient of determination (R^2)). Therefore, the probability that a user will want to watch a movie as a function of its rank (to the specific user) takes the form of (where α and β are constants):

$$p(m|r(m)) = \alpha - \beta \cdot r(m) \quad (1)$$

Given a new user, PUMA sorts the list of movies which is outputted by the original recommender system according to its expected promotion value, which is given by:

$$p(m|r(m)) \cdot v(m) \quad (2)$$

and provides the top n movies as its recommendation.

3.2 Algorithm for Revenue Maximizing

In this setting, every movie is assigned a fixed price (different movies have different prices). Each movie is also assumed to have a cost to the vendor. PUMA intends to maximize the revenue obtained by the vendor, which is the sum of all movies purchased by the users minus the sum of all costs to the vendor.

PUMA’s variant for the Revenue Maximizing settings confronts a much more complex problem than the PUMA’s variant for the hidden agenda for the following two reasons: 1. There is a direct conflict between the system and the users. 2. PUMA must model the likelihood that the users will watch a movie as a function of both the movie rank *and* the movie price.

Building a model by learning a function of both the movie rank and the movie price together is unfeasible as it requires too many data points. Furthermore, in such a learning phase the movie supplier intentionally provides sub-optimal recommendation, which may result in a great loss. Instead, we assume that the two variables are independent, i.e. if the movie rank drops, the likelihood of the user buying the movie drops similarly for any price group. This assumption was used since we found no correlation between the prices of movies in Amazon and their popularity in IMBD.

In order to learn the impact of the price on the likelihood of the users buying a movie, we use the recommender system as is, providing recommendations from 1 to n . We cluster the data into pricing sets where each price (fee f) is associated with the fraction of users who want to buy a movie (m) for that price. Using least squares regression we find a function that best explains the data as a function of the price. We tested the same functions described above (see Table 1 - replace movie rank with movie fee), and the log function resulted with a nearly perfect fit to the data. Therefore, the probability that a user will be willing to pay in order to watch a movie as a function of its fee takes the form of (where α and β are constants):

$$p(m|f(m)) = \alpha_1 - \beta_1 \cdot \ln(f(m)) \quad (3)$$

In order to learn the impact of the movie rank (r) in the recommender system on the likelihood of the users buying a movie, we removed all prices from the movies and asked the subjects if they were willing to pay to watch the movie (without mentioning its price). As in the hidden agenda settings, we provided recommendations in leaps of k' (i.e. recommendations are in the group $\{1, k' + 1, \dots, (n - 1) \cdot k' + 1\}$). We clustered the data according to the movie rank and once again using least squared regression we found a function that best explains the data as a function of the movie rank. Among the functions that we tested (see Table 1), the log function turned out to provide the best fit to the data for the movie rank as well (resulting with the greatest coefficient of determination (R^2)). Using the log function (which is a convex function) implies that the drop in user acceptance rate between movies in the top rankings is larger than the drop in user acceptance rate within the bottom rankings. The difference in the function which best fits the data between the hidden agenda setting and the revenue maximizing setting is sensible, since, when people must pay for movies they are more keen that the movies be closer to their exact taste, therefore the acceptance rate drops more drastically. The probability that a user will be willing to pay in order to watch a movie as a function of its rank takes the form of:

$$p(m|r(m)) = \alpha_2 - \beta_2 \cdot \ln(r(m)) \quad (4)$$

A human model for predicting the human willingness to pay to watch a movie, $p(m|r(m), f(m))$, requires combining Equations 3 and 4; however this task is non-trivial. Taking $p(m|r(m))f(m)$ to be $p(m|r(m)) \cdot p(m|f(m))$ does not make sense: for example if both signals say that the probability of watching is 0.5 then the output should be 0.5 and not 0.25. Using this intuition, we assume that Equation 3 is exact for the average rank it was trained upon which is $\frac{n}{2} + 1$. Therefore, by adding a correction term, $\gamma(m)$, to Equation 4 we require that Equation 4 gives the same viewing probability as Equation 3 on $\frac{n}{2} + 1$:

$$\alpha_2 + \gamma(f(m)) - \beta_2 \cdot \ln(\frac{n}{2} + 1) = \alpha_1 - \beta_1 \cdot \ln(f(m)) \quad (5)$$

Isolating $\gamma(m)$ we get:

$$\gamma(f(m)) = (\alpha_1 - \alpha_2) + \beta_2 \cdot \ln(\frac{n}{2} + 1) - \beta_1 \cdot \ln(f(m)) \quad (6)$$

Therefore, our human model for predicting the fraction of users who will buy a movie, m , given the movie price, $f(m)$, and the movie rank, $r(m)$ (obtained from the recommender system) is:

$$p(m|r(m), f(m)) = \alpha_2 + ((\alpha_1 - \alpha_2) + \beta_2 \cdot \ln(\frac{n}{2} + 1) - \beta_1 \cdot \ln(f(m))) - \beta_2 \cdot \ln(r(m)) \quad (7)$$

and after simple mathematical manipulations:

$$p(m|r(m), f(m)) = \alpha_1 - \beta_2 \cdot \ln(\frac{r(m)}{\frac{n}{2} + 1}) - \beta_1 \cdot \ln(f(m))$$

Once a human model is obtained, PUMA calculates the expected revenue from each movie simply by multiplying the movie revenue with the probability that the user will be willing to pay to watch it (obtained from the model) and returns the movies with the highest expected revenues. The revenue is simply the movie price ($f(m)$)



Figure 1: A screen-shot of a subject selecting movies he liked

minus the movie cost to the vendor ($c(m)$). I.e. given a human model, PUMA recommends the top n movies which maximize:

$$(f(m) - c(m)) \cdot p(m|r(m), f(m)) \quad (8)$$

4. EXPERIMENTS

All of our experiments were performed using Amazon's Mechanical Turk service (AMT). Participation in all experiments consisted of a total of 245 subjects from the USA, of which 50.6% were females and 49.4% were males, with an average age of 31.5. The subjects were paid 25 cents for participating in the study and a bonus of additional 25 cents after completing it. We ensured that every subject would participate only once (even when considering different experiments). The movie corpus included 16,327 movies. The original movie recommender system receives a list of 12 preferred movies for each user and returns a ranked list of movies that have a semantically similar description to the input movies, have a similar genre and also considers the released year and the popularity of the movies (a personalized non-collaborative filtering-based recommender system). We set $n = 10$, i.e., each subject was recommended 10 movies.

After collecting demographic data, the subjects were asked to choose 12 movies which they enjoyed most among a list of 120 popular movies (see Figure 1 for a screen-shot). Then, depending on the experiment, the subjects were divided into different treatment groups and received different recommendations.

The list of recommendations included a description of each of the movies (see Figure 2 for an example). The subjects were shown the price of each movie, when relevant, and then according to their treatment group were asked if they would like to pay in order to watch it, or simply if they would like to watch the movie. In order to assure truthful responses, the subjects were also required to explain their choice ("Please explain why (mandatory)"). After receiving the list of recommendations and specifying for each movie if they would like to buy it (watch it), the subjects were shown another page including the exact same movies. This time they were asked whether they have seen each of the movies ("Did you ever watch movie name?"), whether they think that a given movie is a good recommendation ("Is this a good recommendation?") and rated the full list ("How would you rate the full list of recommendations?") on a scale from 1 to 5. These questions were intentionally asked on a different page in order to avoid framing [25] and to ensure that the users return their true preferences².

²We conducted additional experiments where the subjects were first asked whether they watched each movie and then according

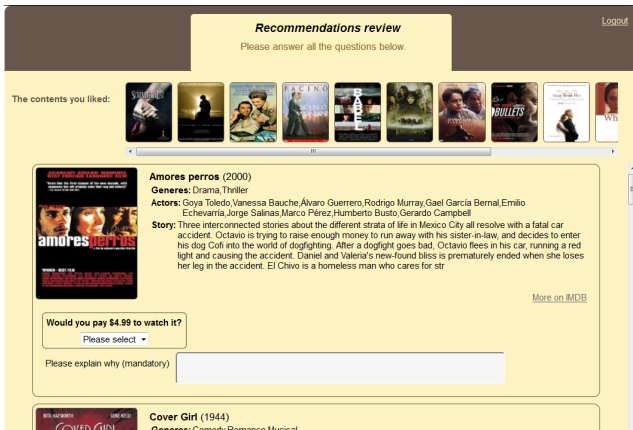


Figure 2: Recommendation page screen-shot

Table 2: Demographic statistics for the hidden agenda setting

| treatment group | number of subjects | fraction of females | average age |
|-----------------|--------------------|---------------------|-------------|
| Rec-HA | 31 | 58.1% | 32.3 |
| PUMA-HA | 30 | 50.0% | 32.6% |
| Learn-HA | 30 | 53.3% | 29.4% |

4.1 Hidden Agenda Setting

In the hidden agenda setting we assume that the subjects have a subscription and therefore they were simply asked if they would like to watch each movie ("Would you watch it?"). The hidden agenda setting experiment was composed of three different treatment groups. Subjects in the *Rec-HA* group received the top 10 movies returned by the original recommender system. Subjects in the *PUMA-HA* group received the movies chosen by PUMA. Subjects in the *Learn-HA* group were used for data collection in order to learn PUMA's human model. Table 2 presents some demographic statistics on the subjects in these three treatment groups.

For the data collection on the movie rank phase (*Learn-HA*) we had to select a value for k (which determines the movie ranks on which we collect data; see Section 3.1). The lower the k is, the more accurate the human model is for the better (lower) rankings. However, on the other hand, the higher k is, the more rankings the human model may cover. In the extreme case where the ranking has a minor effect on the human acceptance rate, the vendor may want to recommend only movies with a promotion value of 1. Even in that extreme case, the highest movie rank, on average, should not exceed $|V| \cdot n$, which is 100. Therefore, we set $k = 10$, which allows us to collect data on movies ranked: $\{1, 11, 21, 31, 41, 51, 61, 71, 81, 91\}$.

Unfortunately, the data collected in the *Learn-HA* group was very noisy, as the movies in the 11th rank resulted with a much higher acceptance rate than those in the 1st rank. Furthermore, the movies in the 71st rank resulted in a much lower acceptance rate than those in the 81st rank. Therefore the coefficient of determination (R^2) was only 0.31. Still, the tendency of the data was clear (the correlation between the movie rank and the acceptance rate

to their answer, were asked whether they would pay for watching it (again). We obtained similar results to those presented in Section 4.3, regarding peoples preference to movies that they have already seen. However, we do not include these results, since they may have been contaminated by the framing effect.

Table 3: Coefficient of determination for functions tested for the hidden agenda setting

| function | R^2 |
|----------|-------------|
| linear | 0.31 |
| exponent | 0.29 |
| log | 0.21 |
| power | 0.21 |

Table 4: The fraction of subjects who wanted to watch each movie, average promotion gain, overall satisfaction and fraction of movies who were marked as good recommendations

| treatment group | want to watch | average promotion | overall satisfaction |
|-----------------|---------------|-------------------|----------------------|
| Rec-HA | 76.8% | 0.436 | 4.13 |
| PUMA-HA | 69.8% | 0.684 | 3.83 |
| Learn-HA | 62.0% | - | 3.77 |

was negative 0.56, which implies that the original recommender system performed well), and additional data would have probably yielded a better coefficient of determination. Nevertheless, the fit-to-data reached was definitely good enough, as can be seen by the performance results. Table 3 lists coefficient of determination for all functions tested.

The specific human model obtained, which was used by PUMA (in the hidden agenda settings) is simply:

$$p(m|r(m)) = 0.6965 - 0.0017 \cdot r(m) \quad (9)$$

As for the results: PUMA significantly ($p < 0.001$ using student t-test) outperformed the original recommender system by increasing its promotion value by 57% with an average of 0.684 per movie for PUMA-HA versus an average of only 0.436 per movie for the *Rec-HA* group. No statistically significant differences were observed between the two groups from the average satisfaction for each of the movies or in the user satisfaction from the full list. It is likely that using more data, one could get a statistically significant drop in user satisfaction. Our best estimate is a 3% drop in the fraction of good recommendations (from 71% rated as good recommendations in the *Rec-HA* group vs. 69% in the *PUMA-HA* group), and a 7% loss in the satisfaction from the entire list. See Table 4 for additional details.

4.2 Revenue Maximizing Setting

For the revenue maximizing settings, all movies were randomly assigned a price which was in $F = \{\$0.99, \$2.99, \$4.99, \$6.99, \$8.99\}$.

³ We assumed that the vendor's cost doesn't depend on the number of movies sold and therefore set $c(m) = 0$ for all movies. The subjects were asked if they would pay the movie price in order to watch the movie ("would you pay \$movie price to watch it?"). As in the hidden agenda setting, subjects were divided into three treatment groups. Subjects in the *Rec-RM* group received the top 10 movies returned by the original recommender system. Subjects in the *PUMA-RM* group received the movies chosen by PUMA. Subjects in the *Learn-RM* group were used in order to obtain data about the decay of interest in movies as a function of the movie rank (as explained in Section 3.2). The subjects in this group were asked if they were willing to pay for a movie, but were not told its

³We used random pricing since we didn't find correlation between Amazon's movie price and features such as popularity of the movie in IMDB, release year, parental rating and country production.




Figure 3: An example of PUMA’s selection process

Table 5: Demographic stats. for the revenue maximizing setting

| treatment group | number of subjects | fraction of females | average age |
|-----------------|--------------------|---------------------|-------------|
| Rec-RM | 31 | 41.9% | 32.1 |
| PUMA-RM | 28 | 67.8% | 29.7 |
| Learn-RM | 30 | 40.0% | 33.9 |

price ("Would you pay to watch it?"). Table 5 presents some demographic statistics on the subjects in these three treatment groups.

In the movie rank learning phase, we set $k' = 5$, i.e., recommendations were in the group $\{1, 6, 11, 16, 21, 26, 31, 36, 41, 46\}$. Once again, this is because $k' \cdot n = |F| \cdot n$ (even if the movie ranking has minor impact on the probability that the user will watch the movie, and therefore PUMA would stick to a certain price; still, on average, it is not likely that PUMA will provide movies which exceed rank $|F| \cdot n$, and therefore no data is needed on those high rankings). The coefficient of determination value using the log function on the learning data was 0.60, as can be seen in Table 6.

The specific human model obtained, which was used by PUMA (in the revenue maximizing settings), is:

$$p(m|r(m), f(m)) = 0.82 - 0.05 \cdot \ln\left(\frac{r(m)}{6}\right) - 0.31 \cdot \ln(f(m))$$

As can be seen in Figure 4, PUMA significantly ($p < 0.05$ using student t-test) outperformed Rec-RM, yielding an average revenue of \$1.71, as opposed to only \$1.33 obtained by Rec-RM. No significance was obtained when testing the overall satisfaction level from the list: 4.13 vs. 4.04 in favor of the Rec-RM group. However, more data could probably give a statistically significant difference. Our best estimate for this loss would be about 2.2%. While the average movie price was also similar in both groups, with an average movie price of \$5.18 for Rec-RM and an average movie price of \$5.27 for PUMA-RM, the standard deviation was quite different: 2.84 for the Rec-RM group, and only 1.95 for PUMA-RM, in

Table 6: Coefficient of determination for the functions tested for the revenue maximizing setting.

| function | R^2 |
|----------|-------------|
| linear | 0.43 |
| exponent | 0.39 |
| log | 0.60 |
| power | 0.54 |

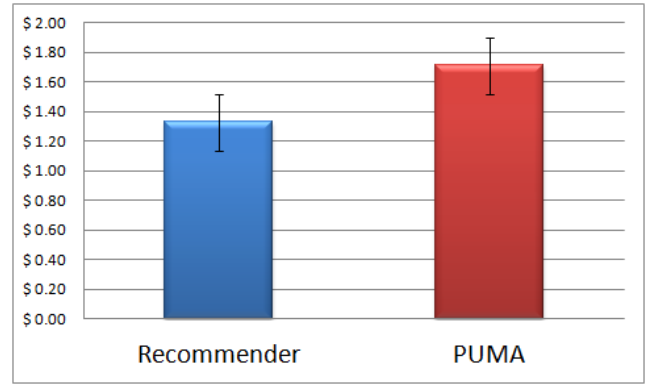


Figure 4: Average revenue for system (in dollars)

Table 7: the fraction of subjects who wanted to pay for a movie, average revenue and overall satisfaction.

| treatment group | want to buy | average revenue | overall satisfaction |
|-----------------|-------------|-----------------|----------------------|
| Rec-RM | 39.1% | \$1.33 | 4.13 |
| PUMA-RM | 37.1% | \$1.71 | 4.04 |
| Learn-RM | 56% | – | 4.03 |

which 64.6% of the movies were either priced at \$2.99 or \$4.99.

Figure 3 demonstrates the selection process performed by PUMA for a specific user. After calculating the expected revenue using the human model, PUMA selects movies #1, #3, #4, #5, #7, #9, #12, #13, #15 and #21, which yield the highest expected profit. In this example, when comparing PUMA’s recommendation’s expected revenue to the expected revenue from the first 10 movies (which would have been selected by the original recommender system), the expected revenue increases from \$1.34 to \$1.52 (13%).

4.3 Subject-Preference for Movies that Have Been Watched Before


We discovered that in all three groups in the revenue maximizing setting, many subjects were willing to pay for movies that they have already watched before. As we discussed in the related work section, marketing literature deals both with the variety effect (buyers who want to enrich their experiences) and with loyalty (or the mere-exposure effect). However, movies are considered to be a prominent example of a variety product, in which customers want to have new experiences, and therefore this result is surprising.

Furthermore, subjects were willing to pay for movies that they have already watched more often than movies which were new to them (although only in the Learn-RM group, where no price was present, these differences reached statistical significance with $p < 0.001$ using Fisher’s exact test). Similar behavior was also observed in the hidden agenda setting, where the subjects in all three groups were simply asked if they would like to watch each movie (and neither the word ‘buy’ nor a price was present). In the hidden agenda settings the subjects significantly ($p < 0.001$) preferred watching a movie again to watching a movie that was new to them. We therefore set out to test whether this behavior will reoccur when the movies are cheap and have a fee of \$0.99, \$1.99 or \$2.99 - the *Cheap* group (see demographic details in Table 8).

This pattern was indeed repeated in the cheap group when prices were mentioned and with statistical significance ($p < 0.001$ using Fisher’s exact test). Figure 5 compares the fraction of subjects who

Table 8: Demographic statistics on the Cheap treatment group

| treatment group | number of subjects | fraction of females | average age |
|-----------------|--------------------|---------------------|-------------|
| Cheap | 35 | 34.3% | 29.3 |

**Figure 5: Comparison between the fraction of subjects who wanted to watch/pay for movies they have already seen and movies that were new to them.**

chose to pay or watch a movie that they hadn't watched before to those who chose to pay or watch a movie that they had watched before (see Table 9 for the exact values). We term this phenomenon the Want To See Again (WATSA) phenomenon. On average 53.8% of the movies recommended were already seen in the past by the subjects. One should not be concerned about the differences in the column heights in Figure 5 across the different treatment groups as, obviously, more subjects wanted to watch the movies for free (in the hidden agenda), where no specific price was mentioned (Learn-RM), and when the movies were cheap than when the movies were more expensive (Rec-RM and PUMA-RM).

In practice, when a recommender system is present, many times the users are instructed to simply choose a (fixed size) set of movies which they would like to watch. We therefore conducted an additional treatment group (*Choose-3*) in which the original experimental settings were slightly modified. In this treatment group the subjects were instructed to choose the movies which they would like to watch ("Assume you are entitled to watch three of the recommended movies listed below, which movies would you choose?"). Please refer to Table 8 for demographic details on the choose-3 treatment group.

As we hypothesized, in the choose-3 treatment group, subjects were significantly ($p < 0.001$) more likely to select movies which they have already seen in the past more than movies which were new to them. While only 8.4% of the movies which were new to the subjects were chosen - 45.0% of the movies that have already been watched were chosen by the subjects. (In total, 59% of the

Table 9: The fraction of subjects who wanted to watch/pay for movies they had already seen and movies that were new to them

| treatment group | seen movies | new movies | delta |
|-----------------|-------------|------------|-------|
| Rec-HA | 92.2% | 55.4% | 36.8% |
| PUMA-HA | 94.7% | 53.7% | 41.0% |
| Learn-HA | 91.1% | 47.2% | 43.9% |
| Rec-RM | 39.1% | 36.6% | 2.5% |
| PUMA-RM | 39.5% | 33.3% | 6.2% |
| Learn-RM | 70.1% | 42.5% | 27.6% |
| Cheap | 49.8% | 28.5% | 21.3% |

Table 10: Demographic statistics on the Choose-3 group

| treatment group | number of subjects | fraction of females | average age |
|-----------------|--------------------|---------------------|-------------|
| Choose-3 | 30 | 60.0% | 33.1 |

movies have already been watched by the subjects.)

All these findings may be very relevant to designers of recommender systems for movies. Today, most systems take great care not to recommend movies that the user has already seen, while instead perhaps one should try to recommend movies that the user has both seen and liked.

5. DISCUSSION AND FUTURE WORK

One may be concerned about PUMA's performance in the long run, when it is required to interact with the same person many times. Although this hasn't been explicitly tested, we believe that PUMA's advantage over the recommendation system will not degrade; as the overall satisfaction from PUMA's recommendations and the average movie fee for PUMA's recommendations (in the revenue maximization setting) are both very close to that of the original recommender system. An interesting property of PUMA is that it allows online learning, as it may collect additional statistics on-the-fly and use it to refine its human model. In the revenue maximization setting, there is a clear conflict between the business and the user: recommending movies that the advertiser prefers (expensive ones) is bound to reduce the probability that suggestions are accepted. In the hidden agenda setting, all movies are a-priori the same for the user, and hence the only loss in showing a recommendation that the business likes to promote is that it's lower on the user's list. Due to this, there is an even greater gain in changing the list of recommendations, and we see a larger gap between PUMA and the recommendation engine in the hidden agenda setting.

The WATSA phenomenon is heavily linked to the mere-exposure effect [26]. It is well known that people prefer to consume products from brands they know, or to consume the same product several times [3]. However, seeing this phenomenon in movies is surprising. Movies are often believed to be novelty driven, and watching the same movie more than once may eliminate many properties upon which movie enjoyment relies heavily (such as the element of surprise, punch lines etc.). We assume that WATSA is also linked to risk aversion, in which people prefer to watch or even pay for movies that they have seen in the past since they have already enjoyed them and are therefore less likely to regret watching them again. This may be a result of people not fully trusting the movie recommender system. A different reason might be that people truly enjoy watching movies that they have seen in the past, more than new movies. This might be due to the fact that most people watch movies that they expect to enjoy, and therefore people are more likely to enjoy movies that they have already seen. An interesting future experiment would be to have subjects watch a set of movies returned by a recommender system. Upon completion, the subjects would be asked how much they enjoyed each movie and then to compare the average enjoyment from already seen movies to new movies. Such an experiment will reveal whether people's preference for watching movies that they have already seen is not justified and may only be superfluous risk aversion, or a true preference that is linked to people's enjoyment from watching movies that have already been seen. Further research is also required here in order to see if there is a difference between movies which were seen recently and ones which were seen a long time ago, how many times a user is likely to want to watch a movie, whether there is a de-

pendency on the genre, culture, etc. It is also very likely that the WATSA phenomenon is a unique property of personalized recommender systems, which supply good recommendations. We leave all these questions for future work.

An interesting property of the WATSA phenomenon may be implied from Table 9: the cheaper the movies are, the greater the WATSA phenomenon. When the prices are the highest (in the Rec-RM and PUMA-RM groups), the difference between the fraction of subjects willing to pay for movies that they have seen and the fraction of subjects willing to pay for new movies is only 2.5%–6.2%. When the movies are cheap, this difference leaps to 21.3%, when no price is mentioned it reaches 27.6%, and when the subjects are just asked whether they would like to watch the movies, this difference shoots up to 36.8%–43.9%! Such behavior may be explained by the fact that people might be willing to pay large amounts only for new movies that they are sure that they would enjoy, and they are willing to pay small amounts for movies that they have enjoyed in the past as they see it as a risk-less investment. However, when testing the prices within the Rec-RM, PUMA-RM and Cheap groups, the WATSA phenomenon clearly increased as the prices decreased only in the PUMA-RM group. In the other two groups (Rec-RM and Cheap), the WATSA phenomenon remained quite steady among the different price groups. This may imply that the average cost has greater impact on the WATSA phenomenon than the specific price of each movie. Therefore this property still requires additional study.

6. CONCLUSIONS

In this paper we introduce PUMA, an algorithm which mounts a given black-boxed movie recommender system and selects movies which it expects will maximize the system’s revenue. PUMA builds a human model which tries to predict the probability that a user will pay for a movie, given its price (if applicable) and its rank in the original recommender system. We consider two different environments, one in which the movies are associated with different prices and the other in which no prices are present. We demonstrate PUMA’s high performance empirically, using an experimental study.

Another important contribution of the paper is the uncovering of a phenomenon in which people prefer watching and even paying for movies which they have already seen (when suggested by a recommender system) to movies which are new to them. This phenomenon was tested and found statistically significant in an extensive experimental study as well.

7. REFERENCES

- [1] A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and Y. Gal. Strategic advice provision in repeated human-agent interactions. In *AAAI*, 2012.
- [2] L. Backstrom and J. Leskovec. Supervised random walks: predicting and recommending links in social networks. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 635–644. ACM, 2011.
- [3] Robert F Bornstein, Paul R D’Agostino, et al. Stimulus recognition and the mere exposure effect. *Journal of personality and social psychology*, 63:545–545, 1992.
- [4] LS. Chen, FH. Hsu, MC. Chen, and YC. Hsu. Developing recommender systems with the consideration of product profitability for sellers. *Information Sciences*, 178(4):1032–1048, 2008.
- [5] A. Das, C. Mathieu, and D. Ricketts. Maximizing profit using recommender systems. *ArXiv e-prints*, August 2009.
- [6] F. Gönül and K. Srinivasan. Modeling multiple sources of heterogeneity in multinomial logit models: Methodological and managerial issues. *Marketing Science*, 12(3):213–229, 1993.
- [7] John R Hauser and Julian Marewski Dzyabura. Consideration-set heuristics. *Journal of Business Research*, (forthcoming), 2010.
- [8] J. Heckman. Heterogeneity and state dependence. In *Studies in labor markets*, pages 91–140. University of Chicago Press, 1981.
- [9] K. Hosanagar, R. Krishnan, and L. Ma. Recommended for you: The impact of profit incentives on the relevance of online recommendations. 2008.
- [10] W. A. Kamakura and G. J. Russell. A probabilistic choice model for market segmentation and elasticity structure. *Journal of Marketing Research*, pages 379–390, 1989.
- [11] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [12] James M Lattin. A model of balanced choice behavior. *Marketing Science*, 6(1):48–65, 1987.
- [13] Andrew Lewis. If you are not paying for it, you’re not the customer; you’re the product being sold. *Metafilter*, 2010.
- [14] Greg Linden, Brent Smith, and Jeremy York. Amazon. com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80, 2003.
- [15] Leigh McAlister and Edgar Pessemier. Variety seeking behavior: An interdisciplinary review. *Journal of Consumer research*, pages 311–322, 1982.
- [16] N. V. Moshkin and R. Shachar. The asymmetric information model of state dependence. *Marketing Science*, 21(4):435–454, 2002.
- [17] Bhavik Pathak, Robert Garfinkel, Ram D Gopal, Rajkumar Venkatesan, and Fang Yin. Empirical analysis of the impact of recommender systems on sales. *Journal of Management Information Systems*, 27(2):159–188, 2010.
- [18] Daniel Read. Monetary incentives, what are they good for? *Journal of Economic Methodology*, 12(2):265–276, 2005.
- [19] F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, editors. *Recommender Systems Handbook*. Springer, 2011.
- [20] Peter E Rossi, Robert E McCulloch, and Greg M Allenby. The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340, 1996.
- [21] C. A. Russell and S. J. Levy. The temporal and focal dynamics of volitional reconsumption: A phenomenological investigation of repeated hedonic experiences. *Journal of Consumer Research*, 39(2):341–359, 2012.
- [22] J Ben Schafer, Joseph Konstan, and John Riedi. Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic commerce*, pages 158–166. ACM, 1999.
- [23] Guy Shani, David Heckerman, and Ronen I. Brafman. An MDP-based recommender system. *J. Mach. Learn. Res.*, 6:1265–1295, 2005.
- [24] W. Shih, S. Kaufman, and D. Spinola. Netflix. *Harvard Business School Case*, 9:607–138, 2007.
- [25] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.
- [26] Robert B Zajonc. Attitudinal effects of mere exposure. 1. *Pers. Soc. Psycho*, 9(2 pt 2), 1968.