

© 2019

Yingqiang Ge

ALL RIGHTS RESERVED

**MAXIMIZING MARGINAL UTILITY PER DOLLAR FOR ECONOMIC
RECOMMENDATION**

By

YINGQIANG GE

A thesis submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Master of Science

Graduate Program in Computer Science

Written under the direction of

Yongfeng Zhang

And approved by

New Brunswick, New Jersey

OCTOBER, 2019

ABSTRACT OF THE THESIS

Maximizing Marginal Utility per Dollar for Economic Recommendation

By YINGQIANG GE

Thesis Director:

Yongfeng Zhang

Understanding the economic nature of consumer decisions in e-Commerce is important to personalized recommendation systems. Established economic theories claim that informed consumers always attempt to maximize their utility by choosing the items of the largest *marginal utility per dollar* (MUD) within their budget. For example, gaining 5 dollars of extra benefit by spending 10 dollars makes a consumer much more satisfied than having the same amount of extra benefit by spending 20 dollars, although the second product may have a higher absolute utility value. Meanwhile, making purchases online may be risky decisions that could cause dissatisfaction. For example, people may give low ratings towards purchased items that they thought they would like when placing the order. Therefore, the design of recommender systems should also take users' risk attitudes into consideration to better learn consumer behaviors.

Motivated by the first consideration, in this paper, we propose a learning algorithm to maximize marginal utility per dollar for recommendation. With the second, economic theory shows that rational people can be arbitrarily close to risk neutral when stakes are arbitrarily small, and this is generally applicable to consumer online purchase behaviors because most people spend a small portion of their total wealth for a single purchase. To integrate this theory with machine learning, we propose to augment MUD optimization with approximate risk-neutral constraint to generate personalized recommendations. Experiments on real-world e-Commerce datasets show that our approach is able to achieve better performance than many classical recommendation methods, in terms of both traditional recommendation measures such as precision and recall, as well as economic measures such as MUD.

Acknowledgements

This research was supported in part by Rutgers University, and I would like to thank my colleagues at Rutgers Computer Science Department who provided insights and expertise that greatly assisted the research.

I would like to thank Professor Yongfeng Zhang for his advising during this research. I also would like to thank Shuyuan Xu for the fruitful discussions, ideas, and his kind assistance in implementation of this work, as well as PhD students Shuchang Liu, Shijie Geng, and Zuohui Fu for their comments that greatly improved the manuscript. Finally, thousands of thanks for my parents, without their mental and physical support, I wouldn't go this far.

Table of Contents

Abstract	ii
Acknowledgements	iii
1. Introduction	1
2. Related Work	4
2.1. Collaborative Filtering	4
2.2. Economic Recommendation	5
3. Preliminaries	6
3.1. Utility and Marginal Utility	6
3.2. Expected Utility and Utility of Expectation	6
3.3. Marginal Utility per Dollar (MUD)	7
3.4. Risk Attitude	8
4. The Framework	10
4.1. The Personalized MUD Function	10
4.2. Multinomial Logistic Modeling (MLM)	11
4.3. Estimated Risk Distribution	11
4.4. Risk-based Optimization of MUD (ROM)	12
4.5. Model Specification	13
4.6. Top-K Recommendation	13
5. Experiments	14
5.1. Dataset Description	14
5.2. Experimental Setup	14

5.3. Evaluation Measure	15
5.4. Experimental Results	15
5.4.1. Recommendation Performance	16
5.4.2. Economic Performance	16
5.4.3. Shallow vs. Deep Models	16
6. Conclusions and Future Work	19
References	20

Chapter 1

Introduction

Because of the significant information overload in the Web environment, designing personalized recommendation systems is important to help users find relevant items efficiently. On one hand, it helps to save the exploration time for consumers, and on the other hand, it helps to improve the revenue of various online economic platforms. Traditional recommendation algorithms mostly focus on optimizing rating- or ranking-oriented measures. For example, rating prediction algorithms such as matrix factorization [10, 12, 23] or neural approaches [27] devise models to optimize for the prediction accuracy in terms of RMSE, while top-N recommendation algorithms such as pair-wise learning to rank [20, 28, 7] propose models to optimize for ranking performance in terms of Precision, Recall, F1 and NDCG, etc.

However, previous recommendation algorithms seldom consider users' economic incentives when modeling the user behaviors and generating recommendations. Actually, consumers' economic incentives play an important role when making decisions in online economic systems such as e-Commerce. By integrating well-established behavioral economic principles and machine learning algorithms, it is possible to develop economics-driven recommender systems so as to make informed recommendations. Therefore, in this paper, we propose to maximize the marginal utility per dollar for economic recommendation. According to behavioral economic theories, rational consumers would always attempt to maximize their utility out of the purchased products by choosing the items that have the largest Marginal Utility per Dollar (MUD) within their budget. The underlying intuition is that consumers would like to spend their money in an efficient way, so that every spent dollar should bring the maximum marginal utility. Meanwhile, making purchasing decisions online involves potential risks of dissatisfaction. A common observation in practical systems is that people may spend money to purchase

an item that they think will match their preference when placing the order, but they may eventually find themselves unsatisfied with the decision when the item is received, and will thus make lower ratings towards the item. As a result, an informed recommendation system should be able to incorporate risk attitude into the model. Fortunately, when integrated with machine learning, established economic principles can help to model the risk attitudes of user decisions based on large-scale user transaction logs.

In this work, we introduce users' risk attitudes into the economic modeling of recommendation systems. In particular, we simulate the risk distribution of each item based on its rating distribution, and by jointly optimizing the marginal utility per dollar under approximate risk neutral constraint, our model learns to predict the consumer decisions with risk-awareness. Experiments based on real-world e-Commerce datasets verify that our approach not only achieves better recommendation performance in terms of precision and recall than both shallow and deep baselines, but also achieves better economic performance in terms of the marginal utility per dollar, which means that our recommendation results would help consumers to spend their money in a more efficient way.

The key contributions of the paper can be summarized as follows:

- We propose to conduct marginal utility per dollar maximization for economic recommendation, which is consistent with established behavioral economic principles, and helps users to spend their money more efficiently.
- We take consumers' risk-attitude into consideration for economic recommendation, which better simulates the real-world online economic environment, where users have to make decisions under potential risks of dissatisfaction.
- We design a joint learning framework to maximize the utility per dollar under risk constraints. Experimental results on several real-world e-commerce datasets show that our approach not only achieves better performance than both shallow and neural recommendation baselines, but also improves the economic measures in terms of money efficiency.

The following parts of the paper will be organized as follows: we review the related work in chapter 2, and then introduce the necessary economics preliminary knowledge in chapter 3

to help readers better understand the background. The proposed model and recommendation strategy are introduced in chapter 4, followed by the experimental results in chapter 5. We finally conclude the work with possible future research directions in chapter 6.

Chapter 2

Related Work

2.1 Collaborative Filtering

Collaborative Filtering (CF) has been an important approach to recommender systems. Early approaches to CF consider the user-item rating matrix and conduct rating prediction with user-based [21, 9] or item-based [24, 14] collaborative filtering methods. With the development of dimension reduction methods, latent factor models such as matrix factorization are later widely adopted in recommender systems, such as singular value decomposition [11], non-negative matrix factorization [13], probabilistic matrix factorization [16], localized matrix factorization [29], etc. In these approaches, each user and item is learned as a latent factor representation to calculate the matching score of the user-item pairs.

Recently, the development of deep learning and neural networks has further extended collaborative filtering methods for recommendation. The relevant methods can be broadly classified into two sub-categories: similarity learning approach, and representation learning approach. The similarity learning approach adopts simple user/item representations (such as one-hot) and learns a complex prediction network as the similarity function to calculate user-item matching scores [8], while the representation learning approach learns rich user/item representations and adopts a simple similarity function (e.g., inner product) for matching score calculation [28].

Another important research direction is learning to rank for recommendation, which learns the relative ordering of items instead of the absolute preference scores. The most representative method on this direction is perhaps Bayesian personalized ranking [20], which is a pair-wise learning to rank method. It is also further generalized to take other information sources such as images [7].

2.2 Economic Recommendation

For a long time, recommendation system research has been working on the above mentioned rating- or ranking-related tasks such as rating prediction and top-N recommendation. However, the related methods seldom consider the economic value that a recommendation list brings to the user or the system, although this is one of the most important goals for real-world recommendation systems. Some recent research on economic recommendation has begun to take care of the economic value of personalized recommendation. For example, [26] studied user's sense of value in terms of utility in recommender systems, and [31] conducted large-scale experiment with real-world users to validate the consumer sense of utility for personalized promotion. [30] further bridged economic principles and machine learning to maximize the social surplus for recommendation, [32] proposed to learn the substitutive and complementary relations between products for utility maximization, and [18] proposed value-aware recommendation for profit maximization based on reinforcement learning. Although current economic recommendation approaches may improve the economic value, their basic motivation is to maximize a total utility function for each user to generate recommendations. However, established behavioral economic principles show that consumers tend to rely on the marginal utility per dollar to make purchasing decisions so as to improve their money efficiency [22, 4], which motivates us to estimate and maximize the marginal utility per dollar for economics-driven recommender systems.

Chapter 3

Preliminaries

3.1 Utility and Marginal Utility

Utility is an economic measure to quantify consumer's pleasure or satisfaction towards some items. Basically, it is widely used to analyze the human behavior in rational choice theory [3]. Utility is usually a function of the consumption quantity Q , which measures the total utility of consuming Q services of a certain product. Among the many different forms of utility functions, the most fundamental and frequently used one is the King-Plosser-Rebelo (KPR) Utility:

$$U_{KPR}(x) = a_x \ln(1 + Q_x), \quad (3.1)$$

where Q_x is the consumption quantity of item x , and a_x is a scale parameter that could be different for different users, which means that people may have different utility even for the same amount of the same product, because of their different personalized preference.

Marginal Utility (MU) is the increment of utility gained from purchasing one extra unit of a product. It can be defined as the first derivative of the utility function, which is,

$$MU_{KPR}(x) = \frac{a_x}{1 + Q_x}. \quad (3.2)$$

3.2 Expected Utility and Utility of Expectation

Expected utility, also known as the Von Neumann-Morgenstern (VNM) utility, is a basic concept in economics, game theory and decision theory to measure user preference in uncertain circumstances. When an individual has to make a decision under uncertainty, it is rational to make a choice with the highest expected utility [25, 5].




Figure 3.1: Representing three different risk attitudes based on the relationship between the expected utility and the utility of expectation. Here, x -axis represents the possible consumption quantity and y -axis is the utility of the corresponding quantity.

Assuming that U is the utility function for random variable X , then the expected utility (EU) is $EU(X) = \sum_{x \in X} U(x)P(x)$. Meanwhile, the utility of expectation is the utility of the expected value of random variable X , which is $U(EX) = U\left(\sum_{x \in X} xP(x)\right)$. In this paper, we use the rating distribution of an item to estimate the probability distribution $P(x)$ for the item.

3.3 Marginal Utility per Dollar (MUD)

In rational choice theory, the condition where a consumer's total utility reaches its maximum is called the *consumer equilibrium*, which means that the consumer would not purchase either more or less products. The equilibrium condition can be defined as follows,

$$\begin{cases} p_1 Q_1 + p_2 Q_2 + \dots + p_n Q_n = I & (3.3) \\ \frac{MU_1}{p_1} = \frac{MU_2}{p_2} = \dots = \frac{MU_n}{p_n} = \lambda & (3.4) \end{cases}$$

where Q_i is the purchasing quantity of product i ; p_i is the price of i ; I is the budget (or 'income' in economics); MU_i is the marginal utility of i ; λ is called the *marginal utility of money*, which is a constant value at the equilibrium [17, 22]. Based on the above KPR utility function, we can also get the functional form of marginal utility per dollar (MUD) as follows,

$$MUD_{KPR}(x) = \frac{a_x}{(1 + Q_x)p_x} \quad (3.5)$$

Intuitively, rational choice and consumer equilibrium means that when deciding which item to buy, users will always choose the one that makes the best use of their money. However, after an item is purchased, its marginal utility (and thus MUD) will diminish, which means that

the next optimal choice could be another item. This process will iterate until equilibrium, where MUDs of all items are at the same level that equals the marginal utility of money λ . More intuitively, at the consumer equilibrium condition, the marginal utility of purchasing any product with the remaining money will be less than the marginal utility of the remaining money itself, so the consumer will stop buying any product. In this work, we use sigmoid-normalized price $\sigma(p_x)$ in Eq.(3.5) for easier optimization.

3.4 Risk Attitude

Risk usually exists when people make decisions under uncertainty. In recommender systems, users never know for sure if they will satisfy with a product or not. Different people have different attitudes towards risk. Let X be a random quantity, whose expected value is $EX = \sum_{x \in X} xp(x)$ (or $EX = \int xp(x)dx$ for continuous cases), then we define *utility of expectation* as $U(EX)$. Generally, risk attitudes can be divided into three categories, depending on the relationship between the expected utility and utility of expectation.

Risk Aversion implies that the expected utility is less than the utility of expectation, as in Figure 3.1(a), $EU(X) < U(EX)$. A risk aversion person always tries to avoid risks.

Risk Appetite means that the expected utility is greater than the utility of expectation, as in Figure 3.1(b), $EU(X) > U(EX)$. A risk appetite person prefers to maximize the profits even with considerable risk.

Risk Neutral implies that the expected utility is equal to the utility of expectation, as in Figure 3.1(c), which means that a person keep neutral in face of risks, $EU(X) = U(EX)$.

For example, suppose there are two choices in a gamble, one is to win either \$100 or \$0 with 50% probability each, and the other choice is to get \$50 with 100% probability. Though the expected utility of both choices are the same (\$50), user's utility of expectation can be different because of their different risk preferences.

In this paper, our model is based on an economic assumption that the risk attitudes of all the consumers is approximately risk-neutral, which means that the optimization algorithm attempts to minimize the difference between $EU(X)$ and $U(EX)$. Theoretical basis of this assumption comes from classical economic theorem on expected utility [19, 2, 1], which shows that a

rational person with a differentiable utility function will always prefer to take a sufficiently small stake in any positive-expected-value bet. In other words, rational persons can be (almost everywhere) arbitrarily close to risk neutral when stakes are sufficiently small. Since most people only spend a small part of their assets purchasing an item online, this theorem is suitable for e-Commerce recommendation scenarios.

Chapter 4

The Framework

4.1 The Personalized MUD Function

Different users may have different marginal utility functions, which represent their personalized preferences towards different items. In order to leverage MUD optimization for personalized recommendation, we need to adapt the (non-personalized) MUD function in Eq.(3.5) to personalized settings. First, the parameter a_x in Eq.(3.5) only depends on the product x instead of a particular user-item pair. To solve the problem, we substitute a_x with a_{ij} , where i indicates the i -th user, and j indicates the j -th item. Besides, an important information in recommendation systems is the user-item rating, which contains rich signals about the consumer preference on targeted items. Usually, higher rating scores for a user-item pair imply higher consumer satisfaction, thus brings higher MUD. As a result, we introduce the rating scores into the MUD function. Finally, the personalized MUD function is defined as:

$$\mathcal{M}_{ij}(\hat{r}_{ij}, Q_{ij}) = \frac{\tanh(\hat{r}_{ij})a_{ij}}{(1 + Q_{ij})\sigma(p_j)} \quad (4.1)$$

where Q_{ij} is the consumption quantity of item j by consumer i ; $\sigma(p_j)$ is the sigmoid price of item j , which can be seen as a normalization of item prices; a_{ij} is the personalized shape parameter that determines the scale of MUD; and $\hat{r}_{ij}$ is the predicted rating score of item j by consumer i . For easy optimization, we adopt hyperbolic tangent $\tanh(\cdot)$ to map the ratings in to the range of (0,1).

There are many different ways to predict the user-item rating scores $\hat{r}_{ij}$, including user/item-based collaborative filtering, matrix factorization, or deep neural networks. For simplicity, we take the most commonly used matrix factorization (MF) approach [11] for rating prediction.

4.2 Multinomial Logistic Modeling (MLM)

Discrete choice problem describes a situation when a consumer chooses an option between two or more discrete alternatives. There are many forms of different discrete choice models and they share some similarities. More formally, suppose that at time point t , consumer i chooses item j over a set of some other alternative products $\Omega_{it}(j)$. We define the total choice set as $\Pi_{it} = \{j, \Omega_{it}(j)\}$ and its k -th element is Π_{it}^k ($\Pi_{it}^1 = j$). The probability that consumer i chooses alternative j is expressed as P_{ij} .

Researchers in economics have utilized Random Utility Models (RUM) to deal with this problem [32]. Different from traditional RUMs, in this paper, we adapt the idea of choosing the alternative item that provides the highest utility into choosing the alternative item with the highest MUD, so as to maximize the utility without violating the budget constraint. In this way, we have:

$$\hat{\mathcal{M}}_i(\Pi_{it}^k) = \mathcal{M}_i(\Pi_{it}^k) + \epsilon_k \quad (4.2)$$

where $\mathcal{M}_i(\Pi_{it}^k)$ represents the true MUD of a product combination, and $\hat{\mathcal{M}}_i(\Pi_{it}^k)$ represent the observed MUD. ϵ_k is a random variable capturing the impact of all unknown factors. Thus, the probability a customer chooses Π_{it}^1 (item j) over other alternatives is:

$$P_{ij}(\hat{\mathcal{M}}_i(\Pi_{it}^1) > \hat{\mathcal{M}}_i(\Pi_{it}^k)) = P_{ij}(\epsilon_k - \epsilon_1 < \mathcal{M}_i(\Pi_{it}^1) - \mathcal{M}_i(\Pi_{it}^k)) \quad (4.3)$$

where $k = 2, \dots, |\Pi_{it}|$. If ϵ_1 and ϵ_k obey an i.i.d. extreme value distribution, then we have,

$$P_{ij}(y_{it} = 1) = P_{ij}(\hat{\mathcal{M}}_i(\Pi_{it}^1) > \hat{\mathcal{M}}_i(\Pi_{it}^k)) = \frac{\exp(\mathcal{M}_i(\Pi_{it}^1))}{\sum_{k=1}^{|\Pi_{it}|} \exp(\mathcal{M}_i(\Pi_{it}^k))} \quad (4.4)$$

where y_{it} is an indication function that,

$$y_{it} = \begin{cases} 1 & \mathcal{M}_i(\Pi_{it}^1) > \mathcal{M}_i(\Pi_{it}^k) \quad \forall k \neq 1 \\ 0 & \text{Otherwise} \end{cases} \quad (4.5)$$

4.3 Estimated Risk Distribution

Shopping online is risky because consumer may be unsatisfactory with the purchased items. In order to better understand consumer behaviors, we take risk into consideration based on classical expected utility theory. To calculate the expected utility of a product, we have to know the probability distribution function $P(x)$ of its utility.

Table 4.1: Parameters setting for each datasets and method where K is the latent factor size, lr is learning rate, $\#i$ is the number of iterations and λ is the regularization parameter.

Dataset	Baby				Electronics				Games			
Methods	K	lr	λ	$\#i$	K	lr	λ	$\#i$	K	lr	λ	$\#i$
CF	50	1e-4	1.5	3	100	1e-4	0.5	2	100	1e-5	1	3
BPR	100	1e-4	1	3	50	1e-4	1	3	100	1e-4	1	3
NCF	100	1e-4	1	5	70	1e-4	1	5	100	1e-4	1	5
MPUM	100	1e-4	1	3	70	5e-4	1	3	100	5e-4	1.5	3
ROM	100	5e-4	1	3	70	1e-4	0.5	3	100	1e-4	1.5	3

In this work, we simulate the distribution function for each specific user-item pair as a normal distribution $\mathcal{N}(\hat{r}_{ij}, \sigma_{MF}^2)$, where $\hat{r}_{ij}$ is the predicted rating score from matrix factorization (MF), and $\sigma_{MF} = \sqrt{\frac{1}{N-1} \sum (r_{ij} - \hat{r}_{ij})^2}$ is the unbiased estimation of prediction RMSE. Since the objective function of MF minimizes RMSE, this distribution naturally applies higher probability to those frequent ratings. As a result, the intuition of the setting is that if the predicted rating $\hat{r}_{ij}$ is close to the true rating r_{ij} (i.e., easy to predict), then we should have a higher probability to observe this rating.

4.4 Risk-based Optimization of MUD (ROM)

We introduce our optimization framework based on approximate risk neutral constraint as follows:

$$\begin{cases} \max \sum_{(i,j) \in \mathcal{R}} \log(P(y_{ij} = 1)) - \lambda \|\Phi\|^2 \end{cases} \quad (4.6)$$

$$\begin{cases} s.t. \min \sum_{(i,j) \in \mathcal{R}} (\bar{U}_{ij} - U_{ij}(\bar{r}))^2 + \lambda \|\Phi\|^2 \end{cases} \quad (4.7)$$

where $\mathcal{R}$ is the set of all the observed user-to-item rating pairs; $P(y_{it} = 1)$ is the multinomial logistic model described in Eq.(4.4); $\bar{U}_{ij}$ is the expected utility, and $U_{ij}(\bar{r})$ is the utility of expectation, which will be introduced in detail in the next section; and finally, Φ is the parameter set to be learned in the corresponding loss function, which will also be crystallized in the next section. The model involves two optimization functions Eq.(4.6) and Eq.(4.7), showing that we maximize the log-likelihood of a multinomial logistic model, so as to optimize the probability of purchasing positive items based on marginal utility per dollar under risk neutral constraint.

4.5 Model Specification

We provide details of the above framework to explain how to learn the parameters. In particular, ratings are predicted based on MF:

$$\hat{r}_{ij} = \alpha^r + \beta_i^r + \gamma_j^r + (\mathbf{p}_i^r)^T \mathbf{q}_j^r \quad (4.8)$$

where α^r is the global bias, β_i^r is the user bias, γ_j^r is the item bias, $\mathbf{p}_i^r$ and $\mathbf{q}_j^r$ are K -dimensional latent factors of user i and item j . Then we use $\hat{r}_{ij}$ as an input to calculate MUD, as shown in Eq.(4.1).

To reduce the parameter space, a_{ij} is also re-parameterized based on the idea of latent factor modeling,

$$a_{ij} = \alpha^a + \beta_i^a + \gamma_j^a + (\mathbf{p}_i^a)^T \mathbf{q}_j^a \quad (4.9)$$

In this way, the optimization problem in Eq.(4.6) and Eq.(4.7) can be specialized as follow:

$$\max \sum_{(i,j) \in \mathcal{R}} \log \left(\frac{\exp(\mathcal{M}_i(\Pi_{it}^1))}{\sum_{k=1}^{|\Pi_{it}|} \exp(\mathcal{M}_i(\Pi_{it}^k))} \right) - \lambda \|\Phi\|^2 \quad (4.10)$$

$$\min \sum_{(i,j) \in \mathcal{R}} \left\{ \sum_{n=1}^5 P_n \tanh(n) - \tanh\left(\sum_{n=1}^5 n P_n\right) \right\}^2 a_{ij}^2 \ln^2(1 + Q_{ij}) + \lambda \|\Phi\|^2 \quad (4.11)$$

where $P_n = Pr(r_{ij} = n) = \int_{n-0.5}^{n+0.5} \mathcal{N}(\hat{r}_{ij}, \sigma_{MF}^2) dr$ is the probability that the rating score of item j is n ($n = 1, 2, 3, 4$ or 5 , representing five different rating scores, for $n = 1$, integral lower bound is $-\infty$, and for $n = 5$, upper bound is $+\infty$). We use SGD optimization algorithm to learn the model parameters.

4.6 Top-K Recommendation

Once we learned the model parameters a_{ij} and $\hat{r}_{ij}$ according to ROM, we can then calculate the marginal utility per dollar $\mathcal{M}_{ij}$ for each user-item pair. We rank all products for a user according to their MUD values, and select the top items whose price is within the user's budget to generate the top-K recommendation list, where in this work, a consumer's budget is considered as the highest price that the consumer had ever spent in his/her purchasing history.

Chapter 5

Experiments

5.1 Dataset Description

We use the consumer transaction data from Amazon¹ [6, 15] in our experiments. The dataset includes user transaction (user id, item id, rating, etc.) and item metadata (item id, price, related item, etc.) on 24 product categories lasting from May 1996 to July 2014. We take three categories (Baby, Electronics and Video Games) that have different size and data sparsity for experiments.

Basic statistics of the experimental datasets as shown in Table 5.1. For each dataset, we sort the transactions of each consumer according to the purchase timestamp, and then split the records into training, validation, and testing sets chronologically by 3:1:1, namely, the first 60% items of each user are used for training, the following 20% for validation, and the last 20% for testing².

5.2 Experimental Setup

We compare our model with the following baselines, including both economic and non-economic methods. For non-economic methods, we involve both shallow and deep learning baselines.

CF: Collaborative Filtering based on matrix factorization is a representative method for rating prediction. In this experiment, we use CF based on latent factor modeling [11].

BPR: Bayesian Personalized Ranking [20] is one of the most representative ranking-based methods for top-N recommendation.

¹<http://jmcauley.ucsd.edu/data/amazon/>

²code of the paper is released at <https://github.com/TobyGE/Maximizing-Marginal-Utility-per-Dollar-for-Economic-Recommendation>




Figure 5.1: Conversion rate on three datasets. x -axis is the length of the recommendation list and y -axis is the conversion rate.

NCF: Neural Collaborative Filtering is a state-of-the-art recommendation algorithm based on deep neural networks. We choose Neural Matrix Factorization to conduct the experiments, which fuses Generalized Matrix Factorization and Multiple Layer Perceptron under the NCF framework.

MPUM: Multi-Product Utility Maximization for recommendation [32], which is an economic recommendation approach that maximizes the utility of product combinations for recommendation.

For each dataset, we use the validation set to find the best parameters of each method, which are shown in Table 4.1.

5.3 Evaluation Measure

To compare the performance of our model and the baselines, we use both traditional measures such as Precision, Recall, and F_1 , and the economic measure MUD. Assuming Γ_i is the recommendation list for consumer i , and N is the total number of consumers, the average marginal utility per dollar over all users is,

$$\mathcal{M}@\Gamma_i = \frac{1}{N} \sum_{i=1}^N \sum_{j \in \Gamma_i} \frac{a_{ij} \tanh(r_{ij})}{(1 + Q_{ij})\sigma(p_j)} \quad (5.1)$$

5.4 Experimental Results

Key experimental results are shown in Table 5.2, and we also plot the conversion rate in Figure 5.1 under different recommendation length K (from 1 to 10). We analyze and discuss the results in terms of the following three perspectives.

5.4.1 Recommendation Performance

According to Table 5.2, among the shallow baseline models (CF, BPR, and MPUM), all pair-wise learning methods (BPR and MPUM) are better than the point-wise CF method, which shows the superiority of pair-wise methods on top-K ranking tasks. By learning deep structures for user-item matching, NCF achieves the best baseline performance in many cases. Furthermore, our ROM approach achieves the best top-K performance against all baselines in most cases. For example, when averaged across all datasets and recommendation lengths, we get 27.97% improvement than NCF. In particular, the improvement is 83.18% for Precision@1 on the *Baby* dataset against NCF. For NDCG, we get 112.54% improvement in average than NCF baseline, especially the largest improvement (407.75%) is achieved on the *Electronics* when $K = 1$. For conversion rate, we get 15.66% improvement than the BPR baseline when averaged across K on the *Baby* dataset, and the largest improvement (63.5%) is achieved when $K = 1$. Compared with NCF, we got 19.01% improvement for top 1 to 10 conversion rate on *Baby* and the largest improvement (78.69%) is achieved when $K = 1$.

The observations imply that by modeling user behaviors based on established economic principles, our model is able to better capture the user preferences for top-K recommendation.

5.4.2 Economic Performance

It is not surprising that our method achieves significantly higher MUD scores than non-economic methods, because we try to optimize for economic values directly. Therefore, we only compare ROM with another economic baseline (i.e., MPUM). We see that our model gets much higher MUD than MPUM, which means that higher utility values does not necessarily imply higher marginal utility per dollar. The higher MUD values produced by our model mean that our recommendations could help users to spend their money more efficiently (i.e., more utility gain per dollar), and thus help to provide more informed recommendations.

5.4.3 Shallow vs. Deep Models

According to Table 5.2 and Figure 5.1, our model is better than NCF in most cases. However, NCF got 6.43% improvement from ROM when $K = 10$ on *Electronics*, and according to

Fig.5.1, NCF grows faster when the recommendation length increases. Neural network is good at learning complex user-item interactions, so it is not surprising that it gets better performance when K is sufficiently large. However, in practical recommender systems the top-5 or 10 results are much more important than later items. Our model is better than NCF from top-1 to top-10 in most cases, and the improvement is especially significant on top-1, which is also reasonable because marginal utility measures the incremental benefit for each extra unit of an item.

Table 5.1: Basic statistics of the experimental data sets.

Dataset	#users	#items	#interactions	sparsity
Baby	23,894	39,767	200,170	0.0211%
Electronics	58,248	45,777	477,074	0.0179%
Video Games	24,735	23,669	236,530	0.0404%

Table 5.2: Summary of the performance. We evaluate for ranking (P, R, F_1 , NDCG) and economic value (MUD), and K is the length of recommendation list. When ROM is the best, its improvements against the best baseline are significant at $p=0.01$.

Dataset	Baby														
Measures	Precision(%)			Recall (%)			F_1 Measure (%)			NDCG			MUD ($\times 10^{-2}$)		
K	1	5	10	1	5	10	1	5	10	1	5	10	1	5	10
CF	3.687	2.658	1.929	2.613	9.421	13.669	3.059	4.147	3.380	0.0084	0.0283	0.0599	0.29448	1.3698	2.6969
NCF	3.829	3.151	2.522	2.714	8.935	17.878	3.177	4.659	4.421	0.0187	0.0503	0.0868	0.28657	1.3699	2.7224
BPR	4.185	3.041	2.483	2.966	10.777	17.600	3.472	4.743	4.352	0.0162	0.0472	0.0837	0.26216	1.3164	2.6720
MPUM	6.177	2.978	2.262	4.378	10.554	16.036	5.125	4.645	3.966	0.0338	0.0718	0.0915	0.26664	1.3482	2.7099
ROM	7.014	3.421	2.553	4.972	12.124	18.098	5.819	5.336	4.475	0.0380	0.0804	0.1021	33.149	73.634	106.523

Dataset	Electronics														
Measures	Precision(%)			Recall(%)			F_1 Measure(%)			NDCG			MUD ($\times 10^{-2}$)		
K	1	5	10	1	5	10	1	5	10	1	5	10	1	5	10
CF	5.983	3.199	2.262	4.316	11.538	16.319	5.014	5.009	3.974	0.0106	0.0336	0.0726	0.01793	0.06293	0.12605
NCF	6.852	4.089	2.944	4.942	14.746	21.233	5.742	6.402	5.170	0.0142	0.0456	0.0970	0.01426	0.05981	0.11355
BPR	7.847	4.011	2.648	5.660	14.468	19.099	6.577	6.281	4.651	0.0087	0.0313	0.0811	0.01149	0.07350	0.12193
MPUM	9.866	3.452	2.333	7.117	12.452	16.827	8.269	5.406	4.098	0.0652	0.1006	0.1168	0.00788	0.05266	0.11403
ROM	11.262	4.159	2.766	8.124	15.000	19.952	9.439	6.512	4.859	0.0721	0.1167	0.1346	14.551	29.770	41.179

Dataset	Video Games														
Measures	Precision(%)			Recall(%)			F_1 Measure(%)			NDCG			MUD ($\times 10^{-2}$)		
K	1	5	10	1	5	10	1	5	10	1	5	10	1	5	10
CF	3.558	2.421	1.891	2.159	7.345	11.473	2.687	3.641	3.247	0.0121	0.0313	0.0568	0.12979	0.64827	1.2875
NCF	3.869	2.907	2.426	2.348	8.819	14.719	2.922	4.372	4.165	0.0170	0.0431	0.0745	0.13018	0.64710	1.2867
BPR	4.253	2.917	2.359	2.581	8.851	14.316	3.212	4.388	4.051	0.0152	0.0404	0.0715	0.12956	0.64319	1.2845
MPUM	6.642	3.007	2.201	4.030	9.123	13.357	5.017	4.523	3.078	0.0389	0.0690	0.0846	0.12303	0.63397	1.2765
ROM	7.026	3.297	2.516	4.263	10.001	15.268	5.307	4.959	4.320	0.0400	0.0737	0.0920	9.0342	20.267	29.874

Chapter 6

Conclusions and Future Work

In this paper, we bridge economic principles and machine learning for personalized recommendation. In particular, we propose to maximize marginal utility per dollar under an approximate risk-neutral constraint for recommendation. Experimental results verified the effectiveness of our model in terms of top-K recommendation and economic values. In the future, we will relax the risk-neutral constraint to other risk attitude functions for more diverse user modeling. We will also consider other economic principles and/or learning methods to benefit recommendation systems both effectively and economically.

References

- [1] K. J. Arrow. Essays in the theory of risk-bearing. Technical report, 1970.
- [2] S. Benartzi and R. H. Thaler. Myopic loss aversion and the equity premium puzzle. *The quarterly journal of Economics*, 110(1):73–92, 1995.
- [3] J. S. Coleman and T. J. Fararo. Rational Choice Theory. *Nueva York: Sage*, 1992.
- [4] D. Friedman and J. Sákovics. Tractable consumer choice. *Theory and Decision*, 79(2):333–358, 2015.
- [5] J. C. Harsanyi. Von neumann-morgenstern utilities, risk taking, and welfare. In *Arrow and the ascent of modern economic theory*, pages 545–558. Springer, 1987.
- [6] R. He and J. McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on World Wide Web*, pages 507–517, 2016.
- [7] R. He and J. McAuley. Vbpr: Visual bayesian personalized ranking from implicit feedback. In *AAAI*, 2016.
- [8] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *WWW*, pages 173–182, 2017.
- [9] J. A. Konstan, B. N. Miller, D. Maltz, J. L. Herlocker, L. R. Gordon, and J. Riedl. GroupLens: applying collaborative filtering to usenet news. *Communications of the ACM*, 40(3):77–87, 1997.
- [10] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8), 2009.
- [11] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- [12] D. D. Lee and H. S. Seung. Algorithms for Non-negative Matrix Factorization. *NIPS*, pages 556–562, 2001.
- [13] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.
- [14] G. Linden, B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, Jan. 2003.
- [15] J. McAuley, C. Targett, Q. Shi, and A. van den Hengel. Image-based recommendations on styles and substitutes. In *SIGIR*. ACM, 2015.

- [16] A. Mnih and R. R. Salakhutdinov. Probabilistic matrix factorization. In *Advances in neural information processing systems*, pages 1257–1264, 2008.
- [17] J. N. Morgan. Can we measure the marginal utility of money? *Econometrica: Journal of the Econometric Society*, pages 129–152, 1945.
- [18] C. Pei, X. Yang, Q. Cui, X. Lin, F. Sun, P. Jiang, W. Ou, and Y. Zhang. Value-aware recommendation based on reinforcement profit maximization. *WWW*, 2019.
- [19] M. Rabin. Risk aversion and expected-utility theory: A calibration theorem. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, pages 241–252. World Scientific, 2013.
- [20] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 452–461. AUAI Press, 2009.
- [21] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. Grouplens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186. ACM, 1994.
- [22] J. Sakovics, D. Friedman, et al. The marginal utility of money: A modern marshallian approach to consumer choice. Technical report, Edinburgh School of Economics, University of Edinburgh, 2011.
- [23] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using markov chain monte carlo. *ICML*, pages 880–887, 2008.
- [24] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295. ACM, 2001.
- [25] P. Wakker and D. Deneffe. Eliciting von neumann-morgenstern utilities when probabilities are distorted or unknown. *Management science*, 42(8):1131–1150, 1996.
- [26] J. Wang and Y. Zhang. Utilizing marginal net utility for recommendation in e-commerce. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 1003–1012. ACM, 2011.
- [27] S. Zhang, L. Yao, and A. Sun. Deep learning based recommender system: A survey and new perspectives. *arXiv preprint arXiv:1707.07435*, 2017.
- [28] Y. Zhang, Q. Ai, X. Chen, and W. B. Croft. Joint representation learning for top-n recommendation with heterogeneous information sources. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1449–1458. ACM, 2017.
- [29] Y. Zhang, M. Zhang, Y. Liu, S. Ma, and S. Feng. Localized matrix factorization for recommendation based on matrix block diagonal forms. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1511–1520. ACM, 2013.
- [30] Y. Zhang, Q. Zhao, Y. Zhang, D. Friedman, M. Zhang, Y. Liu, and S. Ma. Economic recommendation with surplus maximization. In *WWW*, 2016.

- [31] Q. Zhao, Y. Zhang, D. Friedman, and F. Tan. E-commerce recommendation with personalized promotion. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 219–226. ACM, 2015.
- [32] Q. Zhao, Y. Zhang, Y. Zhang, and D. Friedman. Multi-product utility maximization for economic recommendation. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 435–443. ACM, 2017.