Can User Privacy and Recommendation Performance be Preserved Simultaneously?

Tingting Feng, Yuchun Guo, Yishuai Chen

School of Electrical and Information Engineering

Beijing Jiaotong University, Beijing, China

{11111022, ychguo, yschen}@bjtu.edu.cn

Abstract

In online systems of videos, music or books, users' behaviors are disclosed to the recommender systems to learn their interests. Such a disclosure raises a serious concern in the public for the leak of users' privacy. Meanwhile, some algorithms are proposed to obfuscate users' historical behavior records to protect users' privacy, at the cost of degradation of recommendation accuracy. It is a common belief that such tradeoff is inevitable. In this paper, however, we break this pessimistic belief based on the fact that people's interests are not necessarily limited to items which are geared to a certain gender, age, or profession. Based on this idea, we propose a recommendation-friendly privacy-preserving framework by introducing a privacy-preserving module between a recommender system and user side. For instance, to obfuscate a female user's gender information, the privacy-preserving module adds a set of extra factitious ratings of movies not watched by the given user. These added movies are selected to be those mostly watched by male viewers but interesting the given female user. Extensive experiments show that our algorithm obfuscates users' privacy information, e.g., gender, efficiently, but also maintains or even improves recommendation accuracy.

Keywords: Privacy Protection; Balance Precision; Similarity Metric; Recommendation.

1. Introduction

Recommendation service plays an important role in online video systems, but the possible disclosure of users' privacy via recommendation raises serious concern in the public. The accurate recommendation is of cardinal significance for both the service providers and users. It helps users get satisfying products with ease and in turn increase the revenue for the providers [1]. Hence Netflix ever set \$1M Prize for 10% improvement of recommendation accuracy [2]. However, since a recommender system (RS) makes recommendations via mining user historical behavior to find their interests, it has been found that user personal information such as gender and age can be inferred as well. For instance, one's gender can be inferred from his or her records of watching videos [3, 4], and one's age can be inferred from his or her web browsing history [5]. It worries some users that leak of personal information may lead to malicious attacks, e.g., prank calls, annoying advertisement. As a result, some users are even disgusted with personalized recommendation [6]. It is a burning issue of finding a solution to the dilemma of efficient recommendation and the protection of user identity information [7].

Many existing schemes have protected user privacy successfully during the recommendation processes. And they commonly solve this problem via sacrificing the recommendation accuracy [3, 6, 8]. In contrast, in this paper, we set up a simple and effective framework to protect user privacy information from being inferred, and retain the recommendation accuracy without degradation meanwhile. The key to the success of this framework is a privacy-preserving (PP) module that makes the user privacy information blurred but the interest pattern undistorted and even more distinct towards RS. Since both the advertisers and malicious attackers can divide users into two groups, i.e., target group and non-target group, we exam PP module via binary classification of user privacy in detail.

We illustrate its functionality with a simple example. In an online book system, to avoid old people being troubled with too many elderly-oriented advertisements, the PP module will add a set of extra factitious ratings of unread books which satisfy the following conditions: 1) their readers are mostly youth; 2) they are similar to the books that the user is interested in. Then, we set the ratings of selected books equal to the average ratings given by the youth. With the help of PP module, we can not only protect old people's age from being inferred but also provide them with accurate recommendations simultaneously. In Section 5, we introduce the work steps of PP module in detail.

We take gender preservation as an instance to elaborate our algorithm as does the existing work on privacy preservation in RS [3]. Gender is one of the most basic personal information cared by not only users themselves but also advertisers or other third parties who want to make their commercials more targeted. We evaluate our framework

on the MovieLens and Flixster datasets, which are commonly used in studies on RSs and privacy information obfuscation algorithms [3, 9, 10, 11, 12]. In particular, we use *balanced precision* [4] as the metric to evaluate the accuracy of gender inference to handle the serious imbalance between different user groups in this dataset, e.g., the number of male users is much greater than that of female ones. Such an imbalance is common in real systems [5]. And we also make a complexity analysis of our method in Section 6. The main contributions of this paper are as follows.

- We treat the recommendation-friendly privacy protection (RFPP) problem as a multi-objective optimization problem and propose an efficient heuristic algorithm to solve it. As a result, we can preserve privacy without loss of recommendation accuracy simultaneously.
- We introduce a PP module, which serves between users and RSs to disguise user ratings before they are disclosed to the RS and the privacy inferrer (PI). With this module, the privacy information is blurred but the interest preference of user keeps undistorted.
- We propose a new similarity metric of movies based on the average difference of ratings and the number of users. Evaluation results show that our metric outperforms the existing ones.
- We add factitious ratings based on her/his interests and the average ratings from the opposite privacy group. Our scheme performs better both in the recommendation accuracy and the user privacy protection than the existing works.

The rest of this paper is organized as follows. The related work is presented in Section 2. In Section 3, we describe the RFPP problem, and in Section 4, we give out the prerequisite knowledge. Section 5 introduces the algorithm framework, and Section 6 shows our experiment analysis. Finally, we make a conclusion in Section 7.

2. Related Work

There are mainly three aspects of related work about RS and user privacy.

Recommendation enhancement with extra information. The social relationship or other related information of users increasingly received research interests to be used to improve recommendation accuracy [12, 13, 14]. For example, Yang et al. studied Top-k recommendation problem using social network data from Epinions and Flixster [12]. Observing that users may have different opinions on their shared interests, Gurini et al. presented a sentiment-based approach to Twitter user recommendation [15]. For comparison, they all utilize extra data resources to improve RSs' performance. However our method does not need extra data like social relationship and sentiment information.

Protection of data privacy. How to protect user's privacy, especially in personalized services becomes a research focus in recent years. To decrease users' psychological burden, the work of privacy-preserving collaborative filtering (PPCF) begins with Canny's paper [16], and many related works based on PPCF have been proposed, such as [17]. Moreover, some works utilized randomized perturbation techniques (RPT) to realize PPCF and make data private or avoid greatly exposing users' privacy [18, 19, 20, 21]. Others make use of the clustering-based PPCF [22] or the semi-trusted third party [23, 24]. Furthermore, since anonymity is generally conceived to be an integral part of user's right to privacy, Kambourakis contributes to acquiring a comprehensive view of the anonymity research via examining how anonymity is put to work in practice [25]. These works mainly focus on protection of individual user's interests pattern, and in this paper, we consider the protection of demographic information of individual users.

Protection of demographic information. Although different users may have different privacy concerns, personal demographic information is most private. Hence it is taken into consideration in research of privacy protection. Berkovsky et al. described a privacy-enhanced collaborative filtering by distributing the stored profiles to several repositories [6]. They also proposed an algorithm, which blocks the properties of a single user, but not the aggregated preferences of a large group of users, to be inferred. This algorithm achieved this at the cost of some loss of recommendation accuracy [8]. It brings about a concern that when most of the users in a group have the same demographics like gender, once some users have their gender disclosed, the whole group may suffer the risk of disclosure. Besides, Weinsberg et al. blur user gender via adding extra factitious ratings to users' existing records but also with some loss of recommendation accuracy [3]. To our knowledge, none of these schemes can improve the recommendation accuracy when protect user demographic information. And our early work only protects user

gender information for video systems [26]. The recommendation-friendly obfuscation method proposed in this paper is the first general one to solve this dilemma by only obscuring users' privacy information but not their interest information via adding factitious data for a number of scenarios.

3. Problem Description

Is it possible to preserve privacy without degrading recommendation quality? We regard this problem as RFPP problem, and first give the definition of it. In this section, we also make an analysis of the RFPP problem.

3.1 Problem Definition

The problem studied in this paper is how to design a PP module to minimize the loss of recommendation accuracy and privacy disclosure risk for users, given user behavior records. Specifically, the PP module functions as a filter between user side and RS and the possible PI, as shown in Figure 1. A transformation is applied to the user behavior data feeding into the PP module, which obfuscates the demographic information without distorting the user's interest information. RS and PI both work on the data outgoing from PP module. In the reverse direction, recommendation to the users from the RS is rectified in PP module.

The original PI is chosen according to the type of the protecting privacy. It can be trained as a multi-classifier or a linear classifier for the demographic information like race or age. Since almost all advertisers and malicious have a defined target user group, we take the binary classifier as PI to divide users into target group and non-target group. Since the gender prediction is typical binary classification problem, we take gender as the protection privacy to introduce the function of PP module in the following subsections.

Moreover, the framework can be applied in a number of scenarios. Though we take movie systems to exam the function of our method, and the related systems with items like music, CDs, and books also fit the framework in Figure 1. Once they have recommendation requirement, they need to improve recommendation and protecting user privacy simultaneously. Moreover, browsing, listening, watching, rating, etc. such behavior records can be utilized as same as watching records. Thus, the generality of our method is obvious, and we take the movie systems to show the effectiveness of our framework in this paper.

Fig. 1. Illustration of problem framework.

3.2 Problem Analysis

Denote the loss function of RFPP problem, where R denotes the loss of recommendation accuracy, and P denotes the degree of privacy leakage $L(R, P)$ is a two-dimensional evaluation metric. To minimize $L(R, P)$ is equivalent to improve the recommendation accuracy and protect user privacy information simultaneously. Rather than putting forward a compromise method, we intend to design the PP module to solve this problem effectively and efficiently via adding factitious ratings according to the input data, i.e., users' existing ratings.

So far it is believed that obfuscating the demographic information will decrease the accuracy of recommendation, since there is a dependency between user interests and demographic types. However, the relationship between them is not always absolute due to the complexity of the personality [27]. For example, a female may be also interested in some movies mainly favoured by males, if we can estimate her ratings on such movies that she has not watched, and add such estimations to her existing behavior records in PP module, then her gender can be blurred but her interests will not be distorted. It is vice versa for a male.

3.3 Evaluation Metrics

To evaluate the effect of the PP module, we further develop $L(R, P)$ as (1).

$$\bar{L} = \left(RMSE(\tilde{B}), \left| Pre(\tilde{A}) - Pre_0 \right| \right) \quad (1)$$

where $\tilde{B}$ is the recommendation matrix of user ratings, and $\tilde{A}$ is the obscured rating matrix. $Pre(\tilde{A})$ indicates the accuracy of privacy inference and Pre_0 is the ideal accuracy aiming at the privacy protection.

For the recommendation part, we take root-mean-square error ($RMSE$) to measure the differences between values predicted by the scheme in this paper and the values actually observed. The value of $RMSE$ is explicitly the smaller the better. Since our main purpose is to obfuscate users' privacy, the accuracy of gender inference is most concerned. In [3] it is taken for granted that lower accuracy is better for the sense of gender protection, which is not true in some cases, especially when partial knowledge of user privacy is disclosed. Supposing that the gender obfuscation achieves a very low accuracy, says 0.1, if some users' actual gender were disclosed, other users' gender information would be inferred directly with the inverse of the original output by the gender inferrer. The accuracy would be 0.9 exactly. Hence, we argue that for the proper gender obfuscation, the optimal targeted accuracy of gender inference should be 0.5, i.e., the average accuracy of the wild-guesses. As a result, we update (1) to get the final loss function as (2).

$$\bar{L} = \left(RMSE(\tilde{B}), \left| BPre(\tilde{A}) - 0.5 \right| \right) \quad (2)$$

where $BPre(\tilde{A})$ indicates the *balance precision* of the user privacy inference and 0.5 is the average accuracy of the wild-guesses.

Moreover, for the privacy obfuscation part, we take the *balanced precision* [4] as a metric to evaluate the efficiency of our obfuscation algorithm as shown above. For an imbalance system, e.g., the number of males is much larger than that of females, most users' gender can be disclosed easily just by inferring all users' gender to the gender of the majority group. Since the commonly used metric is unfair to the minority for the gender inferrer, we take the *balanced precision* as accuracy, which we first proposed in [4] to evaluate the effect of gender prediction. To deal with the significant shortcomings of accuracy in the imbalance system such as MovieLens, some balanced measures have been proposed [28, 29]. And Brodersen et al. proposed a definition as the weighted sum of precision for each class, i.e., $\tilde{\rho} = c\rho_p + (1-c)\rho_N$ [30], where ρ_p and ρ_N indicate the precision of the positive and negative group, respectively, parameter $c \in (0,1)$ is the cost associated with the misclassification of a case. To strengthen the protection to the minority user, we determine the cost parameter c in analog with the information conveyed in a communication system, based on the information theory [31]. We propose to take the normalized information as the cost with a misclassification as follows.

$$c = \frac{-\log_2 p}{-\log_2 p - \log_2 (1-p)} \quad (3)$$

where $p \in (0,1)$ is a prior probability of a sample to be positive.

4. Prerequisite Knowledge

In this section, we show the basic recommendation mechanism, the gender inference module and the basic gender protecting method. To make our algorithm general, we take the commonly used obfuscating method.

4.1 Matrix Factorization

To setup recommendation mechanism, we utilize the matrix factorization technique, which becomes the first choice for implementing collaborative filtering, thanks to its attractive accuracy and scalability. We make use of the

simple matrix factorization SVD module as Kantor et al. introduced in [1] and predict the users' ratings by minimizing the regularized squared error as (4).

$$\min_{b_i, q_i, p_u, (u,i) \in K} (r_{ui} - \mu - b_i - b_u - q_i^T p_u)^2 + \lambda(b_i^2 + b_u^2 + \|q_i\|^2 + \|p_u\|^2) \quad (4)$$

where r_{ui} is the predicted rating of movie i rated by user u . μ is the average rating of the whole dataset. b_i and b_u are the rating differences of movie i and user u , respectively. And $q_i^T p_u$ captures the implicit feedback between movie i and user u . We compute b_i , b_u , q_i , p_u by minimizing the *RMSE* using the gradient descent method.

4.2 Logistic Regression

The gender inference problem is, in essence, a classifying problem that users are assigned into positive group or negative group according to their rating behaviors. Commonly used Logistic regression model is applied to train the classifier. Besides, we denote the female and male users by 0 and 1 respectively. Furthermore, the coefficient ω_i^0 (ω_i^1) learned from this module, which indicates the strength of movie i correlated with the female users or male users, will be used in our PP module in the recommendation framework. And the ratio between training dataset and test dataset is 8:2.

4.3 Behavior Obfuscation

Since deleting movie watching records is almost impractical in most situations and changing ratings is more suspicious, we obfuscate user gender information and strengthen the user interests via the commonly used method, i.e., adding extra data to the original data in the user side as (5).

$$\tilde{A} = A + A' \quad (5)$$

where $A \in \mathbb{Z}^{U \times I}$ is the training matrix built with the behavior records selected randomly, and a_{ui} denotes the rating of movie $i \in [1, \dots, I]$ rated by user $u \in [1, \dots, U]$. $A' \in \mathbb{Z}^{U \times I}$ is the factitious adding data chosen by PP module. $\tilde{A} \in \mathbb{Z}^{U \times I}$ is the obfuscation matrix, and will replace the original behavior matrix A to train the RS to avoid the leakage of user demographic privacy e.g., the RS or advertiser can learn users' gender information from the data used to recommendation. How to design the obscure matrix $\tilde{A} \in \mathbb{Z}^{U \times I}$ is the most difference between our works and the existing ones, and we will introduce how to get the obscure matrix $\tilde{A}$ specifically in the next section.

5. Algorithm Framework

The algorithm designed for the PP module consists of two operators, i.e., the forward operator and the reverse operator. The former obfuscates user data before it is collected by RS and PI, and the latter makes the modifier of RS recommendation before it is delivered to the user side. Moreover, we introduce the following algorithm based on MovieLens dataset.

5.1 The Forward Operator

In this section, the forward operator is described in detail. We introduce how to choose the movies in PP module, how to compute the similarities between a pair of movies, and how to get the extra factitious ratings. We also illustrate the mechanism of the forward module targeting gender protection specifically.

By the name of PP module, we intend to show that it lets through user interest information but little of gender information. For example, for a woman, this module will choose a set of extra movies that she probably likes, and estimate the ratings that she would give to each movie in this set. Moreover, considering gender protection, for a female user, the added movies should be associated with the males. Firstly, we show the rules and process of movie selection as follows. And then we compare our similarity metric with the existing ones. Finally, the value estimation method of adding ratings is given.

5.1.1 Rules for Selecting Extra Movies

In order to obfuscate a user's gender information but not her/his interests, we suggest the following 3 rules.

Rule 1. Movies dominant by the opposite gender group. Gender protection is primary target in the design of this module. A dominant user group can be associated to a movie in terms of the female (male) dominancy. The female (male) dominancy of a movie can be obtained with the Logistic regression model learned from the training dataset. For a female user, the extra movies should be selected from the male dominant movies, and vice versa.

Rule 2. Movies interesting to the given user. Good recommendation is the secondary target in the design of the obfuscation mechanism. In order to minimize the *RMSE* of the obscure matrix $\tilde{A}$, i.e., $RMSE(\tilde{A})$ described in Section 6.2, the PP module should emphasize users' personal interests by adding extra movies that are interesting to the given users. It is commonly assumed that a movie similar to the movies that a user likes will be interesting to the user as well.

Rule 3. Movies more popular than others. It is commonly accepted that the preference to popular items benefits the grand performance of a recommendation system. Hence selecting extra movies of higher popularity is the third rule.

5.1.2 Extra Movies Selection Process

We take female as positive group, and male as negative group. The process of movie selection is as follows:

(i) First of all, following *Rule 1*, the top- k positive-dominant movie list and the negative-dominant list are built respectively. In this paper, we choose $k = 1000$.

(ii) Next, for a user u belonging to the positive group, the PP module prepares the set of extra movies to be added according to *Rule 2* and *Rule 3*.

(iii) After that, the PP module controls the number of added movies with their popularity. The more popular movies will be selected.

There are two points to be noticed during the selection process of the extra movies.

(i) For each movie i in user u 's existing records, the PP module traverses the opposite movies' list, and picks up movies which are similar to movie i with a similarity more than β . We adjust β large enough to ensure that the added movies are interesting to user u . By contrast, if a movie whose similarity values are less than β is added, it will weaken u 's personal interests, and seem like noise. Via the experiment of MovieLens and Flixster datasets, we find that for a movie system, the similarity threshold value $\beta=0.95$ has good performance towards gender protection.

(ii) The more popular the added movies are, the better the recommendation accuracy is.

In addition, in our experiment, 80% of records from MovieLens are selected randomly as the training dataset, and the other records are regarded as user behavior in the future. All evaluation results are averaged across 10-fold cross-validation in this paper.

5.1.3 A New Similarity Metric

As common recognition, users usually give similar ratings to movies that are similar to the movies they are interested in. Since movie similarity metric is relevant to the accuracy of recommendation, a suitable similarity metric can help PP module pick the similar movies to be added for each user. The existing metrics, such as *Pearson correlation*, *cosine similarity*, and *adjusted cosine similarity*, mainly measure the rating correlation of two movies, and *Jaccard index* considers the coincidence of users instead of the ratings for two movies, which is somewhat related to the movie popularity. Since for a pair of movies the existing algorithms cannot consider the ratings similarities and the number of their identical users simultaneously, we propose a new calculation method of similarity as follow.

$$sim = 1 - \frac{\sum_{u \in S_{ij}} |r_{ui} - r_{uj}|}{N_{ij}} e^{-N_{ij}}, \quad (6)$$

where S_{ij} is a collection of users who watches both movie i and movie j , $N_{ij} = |S_{ij}|$ is the number of users in S_{ij} . r_{ui}, r_{uj} are the ratings of movie i and movie j , given by user u . The similarity contains two sides, i.e., movies' ratings and popularities. $\sum_{u \in S_{ij}} |r_{ui} - r_{uj}| / N_{ij}$ is the average difference of ratings between movie i and movie j . $e^{-N_{ij}}$ is an adjusted quotient (AQ) to control the effect of the movie popularity. For example, when two pairs of movies have the same value of rating side $\sum_{u \in S_{ij}} |r_{ui} - r_{uj}| / N_{ij}$, the pair attracting more users will have a larger *compound similarity* value caused by a smaller $e^{-N_{ij}}$.

To keep the users' interests undistorted, we compare our calculation with several existing similarities on the MovieLens dataset in Table 1 based on the loss function $\bar{L}$, and define the change of *RMSE* of recommendation as ΔR , and the *balanced precision* of gender inference as $BPre$. In addition the loss function $\bar{L}$, *RMSE*, and $|BPre - 0.5|$ are described specifically in Section 6.2.

As Table 1 lists, the PP module can minimize the metric *BPre* close to the average accuracy of wild-guesses based on any kind of similarity. It means our method can obviously protect users' privacy. Comparing with the existing similarities, only our calculation method of similarity can improve the recommendation with 0.002086 decrease of *RMSE*. Moreover, to build the recommendation-friendly obfuscation algorithm, we should choose a most advantage similarity metric to the RS with an acceptable demographic inference result. Hence, the *compound similarity* is the best and we find that in choosing the extra movie set a user might be interested, the movie popularity and the rating should be covered together, and our method has done.

Table 1. Performance of different similarities calculation methods.

Similarity	$\bar{L}$	
	ΔR	$ BPre-0.5 $
Cosine-based similarity	0.002291	0.015463
Correlation-based similarity	0.005639	0.017560
Adjusted cosine similarity	0.005189	0.024379
Jaccard similarity	0.000075	0.053854
Compound similarity	-0.002086	0.049656

To further verify that our calculation method of similarity is reasonable, we compare various forms of adjusted quotients as listed in Table 2. Obviously, $e^{-N_{ij}}$ is the optimal adjusted quotient. $e^{-N_{ij}}$ not only improves the recommendation performance most, but also has a *BPre* mostly close to the average accuracy of the wild-guesses, i.e., 0.5.

Table 2. The calculation of the similarity between movies.

AQ	$\bar{L}$	
	ΔR	$ BPre-0.5 $
$e^{-N_{ij}}$	-0.002086	0.049656
$e^{-N_{ij}^2}$	-0.000355	0.050329
$1 / N_{ij}$	0.000136	0.055043
$1 / N_{ij}^2$	0.000102	0.054037
1	0.000025	0.052867

5.1.4 Factitious Ratings

It is not occasional or intuitive to take the average rating associated with the opposite gender as the extra factitious rating. For the extra movie j selected for user u , several estimations of rating r_{uj} are compared, including the average rating of the same gender raters, average rating of the opposite gender raters and average rating of all raters.

As Table 3 lists, the opposite gender rating performs best in terms of $\Delta RMSE$ and $|BPre - 0.5|$ both. For a female user u , according to the *Rule 1*, the extra movie j is male dominant, hence the most appropriate estimating for the rating of user u for movie j , is the average rating given by the male raters. And the average rating given by all raters comes second. Besides, even the average rating given by female raters only brings an insignificant increase of $RMSE$.

Table 3. The comparison of various extra factitious ratings.

Factitious rating	$\bar{L}$	
	ΔR	$ BPre - 0.5 $
average rating of all raters of j	-0.000334	0.054783
average rating of same gender raters of j	0.000051	0.054293
average rating of opposite gender raters of j	-0.002086	0.049656

5.2 The Reverse Operator

On the reverse direction, the reverse operator will rectify the recommendation output of the RS input into the PP module before it is delivered to the user side. Since the RS considers the nonzero elements in A' as users' history, it will not recommend the corresponding items to the associated users. To do this, we get the final recommendation matrix based on A' , which is reserved by the PP module, as (7).

$$\tilde{B} = \mathcal{P}(B) \quad (7)$$

where $B \in \mathbb{Z}^{U \times I}$ indicates the original recommendation matrix, b_{ui} denotes the predicted value of item $i \in [1, \dots, I]$ for the user $u \in [1, \dots, U]$ by RS. $\tilde{B}$ is the final recommendation matrix, and $\tilde{b}_{ui}$ is the final modified value given to the associated user by PP module via $\mathcal{P}(\cdot)$ function that appends the records of A' to B . In addition, the recommendation accuracy is tested with the users' future behavior, herein the test dataset. And the protection of user demographic is tested based on the obscured rating matrix $\tilde{A}$.

6. Experiment Analysis

In this section, we introduce the datasets firstly. And then the experiment process will be described in detail. Finally, we evaluate the experiment based on a proper range of extra factitious ratings from three respects, respectively, i.e., privacy performance, recommendation performance and complexity analysis. In addition, we denote p_0 the percentage of added ratings relative to total number of records in the training dataset.

6.1 Data Description

We describe the MovieLens and Flixster datasets in this subsection, and also give the basic statistical analysis of both two datasets.

6.1.1 MovieLens

MovieLens is the frequently used RS, which recommends films for its users to watch, based on their film preferences. It is also a virtual community website, which dataset is from a project of GroupLens [32] Research about personalized RSs. This dataset contains 1,000,209 anonymous ratings of approximately 3,900 movies given by 6,040 users who joined MovieLens in 2000. MovieLens is a gender imbalance system. Only 28.29% of its users are female, and 71.71% of its users are male.

Although according to the World Health Organization standard users of MovieLens system can be divided into three intervals according to their ages (18~44), i.e. Teenage (<18), Youth (18~35), Middle aged and elderly people (>35), we divided users into two groups with the view of economic benefit to test the generality of PP module. Since the youth is always with widest interests in the varied kinds of movies as well as the fashionable products displayed by advertisement to cater for any age interval, both the RSs and the commercial advertisers would like most to know and attract them. Thus, the age of the youth is emergent to be protected. Therefore, we take two age groups, i.e. the youth group and the old and kid group. Obviously, the final distribution of age groups is disproportion. The youth group occupies 72.71% of the total users, and the old and kid group only accounts for 27.29% of the total users.

6.1.2 Flixster

Flixster is an American social movie site for discovering new movies, learning about movies, and meeting others with similar tastes in movies. For further comparison, we choose users with valid gender information, landing between Oct. 1st and Dec. 31st 2009, and their corresponding records of viewer id, video rating during this period. The final dataset contains 994,448 anonymous ratings of 3,796 movies given by 6,843 users, each movie with at least five viewers to decrease the sparseness of rating matrix. In addition, 54.95% of the total users are female, and 45.05% of the total users are male.

6.2 Experiment Process

We show the experiment process of our method in this subsection specifically according to the framework represented in Figure 1.

We imply experiments based on gender and age privacy information respectively. To test the effects of adding ratings, we choose different number of added records based on the popularity rank of videos. The experiment involves the following steps.

(i) Select 80% of user behavior records as training dataset, and obscure user data with selected records as Section 5.1 described. Specifically, for each user, we traverse the opposite-dominant movies' list, and picks up movies which are similar to her/his existing records with a similarity more than β and also in top- k popularity movie set of training dataset. After that, we add factitious ratings of selected movies for this user. Hence, we call the total adding records as rank- k video ratings for simplicity.

(ii) Test the recommendation result based on obscuring data with the rest of original user behavior records.

(iii) Test the privacy protection result based on the obscuring data with the PI training on the original user behavior records.

We show the performance of privacy preservation and that of recommendation with adding factitious rank- k video ratings preferred by the opposite gender (age) class.

6.3 Experiment Evaluation

To evaluate the performance of PP module, we analyze the results with respect to privacy performance, recommendation performance and computing complexity. Figure 2 shows the evaluation results based on MovieLens and Flixster datasets. In Figure 2 (a, b, c), the green square dotted line indicates the value of $|BPre - 0.5|$, and the blue circular dotted line indicates the value of ΔR . The red line is the reference line of $\Delta R = 0$, representing no loss of recommendation accuracy. The left and right vertical axis corresponds to ΔR and $|BPre - 0.5|$ respectively. In Figure 2 (d), the lavender (turquoise) diamond dotted line indicates the distribution of p_0 based on rank- k video ratings with gender information in MovieLens (Flixster) dataset. And the lavender triangle dotted line represents the distribution of p_0 with age information in MovieLens dataset.

6.3.1 Privacy Performance

In this subsection, we show the privacy performance of PP module respectively for gender and age preservation. And we also compare our method with existing similar work. Results verify that PP module can preserve user privacy information successfully.

For gender preservation, we show the results with MovieLens and Flixster datasets respectively. As shown in Figure 2 (a), in MovieLens dataset, when $2 \leq k \leq 10$, $3.89\% \leq p_0 \leq 8.82\%$, $0.013 \leq |BPre - 0.5| < 0.05$ corresponds to the average accuracy of the wild-guesses. And as shown in Figure 2 (b), in Flixster dataset, via adding ratings, $|BPre - 0.5|$ decreases more or less obviously. Most importantly, Weinsberg et al. use 10% extra factitious ratings of MovieLens dataset to obscure user ratings based on their proposed method, and make $|BPre - 0.5| \approx 0.145$, and ΔR is 0.058, While we only add 8.12% ($k = 8$) extra records and make $|BPre - 0.5| \approx 0.003$, and ΔR is only 0.012. It is evident that PP module can make $BPre$ greatly close to the average accuracy of the wild-guesses to protect user gender information effectively.

For age preservation, we apply the heuristic algorithm on the MovieLens dataset. And we show the results based on the loss function $\bar{L}$ in Figure 2 (c) and Table 4 respectively. Although $|BPre - 0.5| \approx 0.1$ is not so perfect, the value nearly reduces three times comparing with no adding. Thus, according the analysis above, PP module can preserve privacy information obviously.

Fig. 2. The value of $\bar{L}$ versus rank- k video ratings. (a, b) Gender-preserved recommendation with MovieLens and Flixster datasets respectively. (c) Age-preserved recommendation with MovieLens dataset. (d) Factitious rating percentage p_0 versus rank- k video ratings with MovieLens and Flixster datasets.

Table 4. The comparison of various extra factitious ratings based on age-preserved recommendation with MovieLens dataset.

p_0	$\bar{L}$	
	ΔR	$ BPre - 0.5 $
0	0	0.270148
0.16%	-0.000214	0.115264
0.25%	0.000051	0.110151
0.31%	-0.000011	0.115264
1.14%	0.001003	0.118513

6.3.2 Recommendation Performance

In this subsection, we show the recommendation performance of PP module, and also compare our method with the realistic values and existing similar work. It has been proved that PP module can preserve accuracy of recommendation effectively.

For gender information, although a small number of extra factitious ratings have very little effect on recommendation accuracy, we can decrease $RMSE$ of commendation practically. As Figure 2 (a, d) shown, in MovieLens dataset, when $k = 2$, $p_0 \approx 0.04\%$, $RMSE$ has the largest decrease. Most importantly, Weinsberg et al. use 10% extra factitious ratings of MovieLens dataset to make ΔR equal to 0.058, $|BPre - 0.5| \approx 0.145$. While we only add 3.89% ($k = 2$) extra records and make ΔR equal to -0.002, $|BPre - 0.5| \approx 0.05$. Furthermore, we have an interesting finding that PP module performs better than works of giving added movie a realistic rating known from the test dataset. In detail, for user u of MovieLens dataset, if movie j to be added belongs to u 's records in the test dataset of RS, we add the realistic rating to u 's existing records, and find ΔR is -0.001 (ours is -0.002). Besides, in Flixster dataset, only adding 0.87% ($k = 1$) ratings of existing records can make $RMSE$ absolutely decreases 0.001.

For age information, adding a very limited number of fictitious ratings brings a rapid decline in the two dimensions of $\tilde{L}$ as shown in Figure 2 (c) and Table 4. It means that no more than 0.2% extra ratings can satisfy the users. Thus, PP module can preserve RS's performance simultaneously.

6.3.3 Complexity Analysis

For the feasibility of implementation in reality, the PP module should obscure user behavior data with ease and bring limit extra fictitious records to RS. We make the complexity analysis of our method as follow.

Actually, online video systems devote to protect privacy of loyal users, who are active and register with personal privacy information. Furthermore thousands of popular videos can cover the most of these users. For Flixster system, the precision difference of gender inference trained on 5,000 viewers and 25,000 viewers is only 0.02, and the former result is nearly 0.74 [3]. For MovieLens dataset, top-500 popular videos can cover all viewers. Thus, PP module can utilize a small dataset to finish most of works offline with ease, such as building top- k videos' lists, computing the similarity of videos, and learning the relationship between p_0 and $RMSE$ ($|BPre - 0.5|$).

Besides, PP module can preserve user privacy information and improve the recommendation quality in an efficient way via adding a quite small amount of extra fictitious records to users' existing behavior records. For RS, as the cost of computing the recommendation list depends on the number of non-zeros (nnz) in SVD module, it requires our method satisfy the following constraint:

$$nnz(A) + nnz(A') \approx nnz(A) \quad (8)$$

where $A \in \mathbb{Z}^{U \times I}$ is the training matrix, and $A' \in \mathbb{Z}^{U \times I}$ is the factitious adding data chosen by PP module. Since the value of $nnz(A)$ is as thousands times as that of $nnz(A')$, $nnz(A')$ can be ignored generally. Thus, the PP module does not increase the complexity of RS significantly in deed.

In addition, for online video systems targeting new videos like movies and TV series which release slowly, PP module can update the work process offline each week. In real applications, PP module can adjust update time to make a tradeoff between accuracy and timeliness of recommendations to cater the demand of various RSs.

7. Conclusion

In this paper, we proposed a general and effective framework to overcome the dilemma between protecting users' privacy and improving recommendation accuracy simultaneously. For the sense of gender protection, we define a perfect protection of gender is a prediction accuracy close to the average accuracy of the wild-guesses, i.e., 0.5. We propose a recommendation-friendly user privacy obfuscation algorithm evaluated by the loss function $L(R, P)$ which considers two metrics, i.e., $RMSE$ and *balanced precision*, and the experiment results show that we can achieve a small decrease of $RMSE$ of recommendation, and make the gender results close to wild-guess. Above

all, we design the selection strategy of choosing extra movies and the method of estimating the factitious ratings, and we propose a new similarity metric. Finally, we make experiment analysis of our algorithm and prove that our algorithm can achieve accurate recommendation and protect users' gender and age at the same time.

Since our algorithm is simple and effective, more research about other privacy information will be considered in the further work via different privacy prediction methods. And we want to experiment on user objective online history like shopping records as well.

Acknowledgements

This work was supported in part by the National Science Foundation of China under Grant No. 61271199, 61301082 and the Fundamental Research Funds for the Central Universities under Grant No. W14JB00500.

References

- [1] Kantor P. B., Ricci F., Rokach L., Shapira B., "Recommender Systems Handbook," *Springer*, 2011.
- [2] NetflixPrize [Online]. Available at <http://www.netflixprize.com/>.
- [3] Weinsberg U., Bhagat S., Ioannidis S., Taft N., "BlurMe: Inferring and Obfuscating User Gender Based on Ratings," In *Proc. of ACM RecSys'12*, 2012.
- [4] Feng T., Guo Y., Chen Y., Tan X., Xu T., Shen B., Zhu W., "Tags and Titles of Videos You Watched Tell Your Gender," ICC'14 CSSMA, 2014.
- [5] Hu J., Zeng H. J., Li H., Niu C., Chen Z., "Demographic Prediction Based on User's Browsing Behavior," In *Proc. of WWW'07*, 2007.
- [6] Berkovsky S., Eytani Y., Kuflik T., Ricci F., "Privacy-Enhanced Collaborative Filtering," *User Modeling Workshop on Privacy-Enhanced Personalization*, 2005.
- [7] Wang Y., Kobsa A., "Privacy-Enhancing Technologies," *Social and Organizational Liabilities in Information Security*, 2006.
- [8] Berkovsky S., Eytani Y., Kuflik T., Ricci F., "Hierarchical Neighborhood Topology for Privacy Enhanced Collaborative Filtering," In *Proc. of CHI PEP'06*, 2006.
- [9] Parameswaran R., Douglas M. B., "Privacy Preserving Collaborative Filtering Using Data Obfuscation," *IEEE GRC'07*, 2007.
- [10] Berkovsky S., Kuflik T., Ricci F., "The Impact of Data Obfuscation on The Accuracy of Collaborative Filtering," *Expert Systems with Applications*, 2012.
- [11] Bilge A., Polat H., "An Improved Profile-Based CF Scheme with Privacy," In *Proc. of IEEE'11*, 2011.
- [12] Yang X., Steck H., Guo Y., Liu Y., "On Top-k Recommendation Using Social Networks," In *Proc. of ACM RecSys'12*, 2012.
- [13] Alexandridis G., Siolas G., Stafylopatis A., "Improving Social Recommendations by Applying A Personalized Item Clustering policy," *RSWeb* 2013.
- [14] Ma X., Wang H., Li H., Liu J., Jiang H., "Enhancing Recommended Video Lists for Youtube-Like Social Media," *IEEE MMSP'12*, 2012.
- [15] Gurini D., Gasparetti F., Micarelli A., Sansonetti G., "A Sentiment-Based Approach to Twitter User Recommendation," *Recsys'13*, 2013.
- [16] Canny J., "Collaborative Filtering with Privacy," In *Proc. of IEEE'02*, 2002.
- [17] Bilge, A., Kaleli, C., Yakut, I., Gunes, I., Polat, H., "A Survey of Privacy-Preserving Collaborative Filtering Schemes," *International Journal of Software Engineering and Knowledge Engineering*, 2013.
- [18] Polat H., Du W., "Privacy-Preserving Collaborative Filtering Using Randomized Perturbation Techniques," In *Proc. of IEEE'03*, 2003.
- [19] Polat H., Du W., "Privacy-Preserving Collaborative Filtering," *Int. J. of Electron. Comm*, 2005.
- [20] Yakut I., Polat H., "Privacy-Preserving Eigentaste-Based Collaborative Filtering," *Advances in Information and Computer Security. Springer Berlin Heidelberg*, 2007.
- [21] Gong S., "Privacy-Preserving Collaborative Filtering Based on Randomized Perturbation Techniques and Secure Multiparty Computation," *International Journal of Advancements in Computing Technology*, 2011.
- [22] Bilge A., Polat H., "A Comparison of Clustering-Based Privacy-Preserving Collaborative Filtering Schemes," *Applied Soft Computing*, 2013.
- [23] Aïmeur E., Brassard G., Fernandez J. M., Onana F. S. M., Rakowski Z., "Experimental Demonstration of A Hybrid Privacy-Preserving Recommender System," *IEEE ARES'08*, 2008.
- [24] Aïmeur E., Brassard G., Fernandez J. M., Onana F. S. M., "Alambic: A Privacy-Preserving Recommender System for Electronic Commerce," *International Journal of Information Security*, 2008.

- [25] Kambourakis G., "Anonymity and Closely Related Terms in The Cyberspace: An Analysis by Example," *Journal of information security and applications*, 2014.
- [26] Feng T., Guo Y., Chen Y., "Can User Gender and Recommendation Performance Be Preserved Simultaneously?" In *Proc. of IEEE ICNC'15*, 2015.
- [27] Bruckman A., "Gender Swapping on The Internet," In *Proc. of INET'93*, 1993.
- [28] Kubat, M., Holte, R., Matwin, S., "Learning When Negative Examples Around," *Lecture Notes in Computer Science*, 1997.
- [29] Sokolova M., Japkowicz N., Szpakowicz S., "Beyond Accuracy, F-Score and ROC: A Family of Discriminant Measures for Performance Evaluation," *Lecture Notes in Computer Science*, 2006.
- [30] Brodersen K. H., Ong C. S., Stephan K. E., Buhmann J. M., "The Balanced Accuracy and Its Posterior Distribution," *ICPR'10*, 2010.
- [31] Shannon C. E., "A Mathematical Theory of Communication," *Bell System Technical Journal*, 1948.
- [32] GroupLens [Online]. Available: <http://www.grouplens.com/node/73>.