

When Differential Privacy Meets Randomized Perturbation: A Hybrid Approach for Privacy-Preserving Recommender System

Xiao Liu¹, An Liu^{1,2}(✉), Xiangliang Zhang², Zhixu Li¹, Guanfeng Liu¹,
Lei Zhao¹, and Xiaofang Zhou³

¹ Soochow University, Suzhou, China
anliu@suda.edu.cn

² King Abdullah University of Science and Technology,
Thuwal, Kingdom of Saudi Arabia

³ University of Queensland, Brisbane, Australia

Abstract. Privacy risks of recommender systems have caused increasing attention. Users' private data is often collected by probably untrusted recommender system in order to provide high-quality recommendation. Meanwhile, malicious attackers may utilize recommendation results to make inferences about other users' private data. Existing approaches focus either on keeping users' private data protected during recommendation computation or on preventing the inference of any single user's data from the recommendation result. However, none is designed for both hiding users' private data and preventing privacy inference. To achieve this goal, we propose in this paper a hybrid approach for privacy-preserving recommender systems by combining differential privacy (DP) with randomized perturbation (RP). We theoretically show the noise added by RP has limited effect on recommendation accuracy and the noise added by DP can be well controlled based on the sensitivity analysis of functions on the perturbed data. Extensive experiments on three large-scale real world datasets show that the hybrid approach generally provides more privacy protection with acceptable recommendation accuracy loss, and surprisingly sometimes achieves better privacy without sacrificing accuracy, thus validating its feasibility in practice.

Keywords: Recommender systems · Privacy-preserving · Differential privacy · Randomized perturbation

1 Introduction

During the last few decades we have witnessed the increasing use of recommender systems in various domains to solve the problem of information seeking in an extremely large volume of content. With the help of recommender systems, customers can quickly find things that are interesting or new by narrowing down the set of choices. Meanwhile, service providers using recommender systems can

increase sales or click-through rate (CTR) by providing personalized service for customers. For example, McKinsey¹ reported that “35% of what consumers purchase on Amazon and 75% of what they watch on Netflix come from product recommendations”.

The benefits brought by recommender systems are significant. However, the use of recommender systems introduces privacy threats and concerns. In order to provide high quality recommendations, recommender systems need to collect customers’ private data, such as history data (e.g., the books bought last month or the movies watched last week) and rating data (e.g., the rate for a book or a movie). However, recommender systems may not be trustable. It is common for customers to raise privacy concerns as the collected data may be shared with, rent or sold to third parties. According to a survey done by PewResearch², “86% of Internet users have taken steps online to remove or mask their digital footprints” and “68% of Internet users believe current laws are not good enough in protecting people’s privacy online”. It is thus crucial to develop technologies that can keep users’ private data protected while enabling personalized recommendation, which is a necessary and beneficial complement to the efforts made in the non-technical domain such as privacy policies and related laws.

1.1 Related Work

Cryptography is one of the most important technologies to realize privacy-preserving recommender systems. Using some well-known encryption algorithms, users can transform their private data from meaningful plaintext to meaningless ciphertext, thus achieving privacy preservation. To enable recommender systems to carry out computation over ciphertext directly, the encryption algorithms to be used have to be homomorphic, that is, the result of operations performed on ciphertext, when decrypted, matches the result of operations performed on the corresponding plaintext. For example, Paillier cryptosystem [14] was employed by Erkin et al. [6] and Ma et al. [10], ElGamal cryptosystem [5] was used by Zhan et al. [17] and Badsha et al. [1], to realize privacy-preserving recommender systems. However, homomorphic encryption is built on expensive public-key cryptography, which is theoretical in nature and cannot be applied in practice due to the prohibitive computation cost. In addition, Nikolaenko et al. [13] and Liu et al. [9] built privacy-preserving recommender systems based on another renowned cryptographic tool, Yao’s garbled circuits [8, 16]. However, these approaches require the existence of a trusted third party, which also hinders their application in practice.

To overcome the weakness of cryptography based techniques, Polat and Du [15] proposed a *Randomized Perturbation* (RP) technique which adds noise to users’ private data before releasing the data to recommender systems. RP is much faster than cryptography based techniques, but this is at the cost of sacrificing recommendation accuracy and privacy protection degree. In particular,

¹ <http://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers>.

² <http://www.pewinternet.org/2013/09/05/anonymity-privacy-and-security-online/>.

the smaller the random noise added, the more accurate the predicted rating. However, the smaller noise also results in weak privacy guarantee. For example, the data reconstruction methods proposed by Zhang et al. [18] can derive more original data when smaller noise is injected. Therefore, a trade-off between accuracy and privacy should be made when applying RP.

The above work aims at keeping users' private data secret during recommendation computation. However, the output of recommender systems can be also utilized by malicious users to make inferences about other users' private data [12], that is, based on the recommendation she gets, a malicious user can guess whether someone else has, for example, bought some book or seen some movie. To avoid this kind of information leakage, *Differential Privacy* (DP) [3, 4] has been introduced into recommender systems recently [7, 11, 12]. By adding noise, DP guarantees the distribution of the recommendation is insensitive to any individual user's data, thus preventing the inference of any single user's data from the recommendation. Due to the injected noise, DP also needs to strike a balance between recommendation accuracy and privacy protection degree. In addition, DP does not protect users' data from recommender systems, as the latter has full access to users' data in the clear.

1.2 Contributions

From the above discussion, it is expected that a privacy-friendly recommender system should respect user privacy at two stages: (1) does not ask users to submit their original data in the data collection stage; and (2) can prevent the inference of any single user's data from the final recommendation result in the normal execution stage. To the best of our knowledge, however, these two aspects have not been considered simultaneously. In this paper, we aim at designing an approach that can hide users' private data and prevent privacy inference simultaneously. At first glance, an intuitive solution is to integrate the techniques mentioned above. Nevertheless, there are some interesting issues worthy of investigation but largely overlooked by recent studies. For example, the amount of noise injected by DP is based on the sensitivity of a query function, which sometimes is not easy to estimate, especially when considering that the underlying data will be disguised by RP or encryption. For another, since DP and RP both introduce noise to original data, can we be certain that the recommendation accuracy will inevitably become worse? Or what is the trade-off between accuracy and privacy in this new context?

As the initial step towards more privacy-friendly recommender systems, we propose in this paper a hybrid approach which combines RP and DP. Specifically, users mask their original data through RP and send the disguised values to the recommender system, which injects calibrated noise again to the perturbed data to achieve DP. Our contributions are summarized as follows:

- We design a hybrid approach for privacy-preserving recommender systems by combining DP with RP. Compared with existing works, our approach provides more privacy guarantee as users' private data is kept secret and no one can infer any single user's data from the recommendation result.

- We theoretically show the noise added by RP has limited effect on recommendation accuracy and the noise added by DP can be well controlled based on the sensitivity analysis of functions on the perturbed data.
- We conduct extensive experiments to evaluate the performance of our hybrid approach on three large-scale real world datasets. The results show that the combination of DP and RP is feasible in practice. Generally it provides more privacy protection with acceptable accuracy loss, and surprisingly sometimes it achieves better privacy and accuracy at the same time.

The rest of the paper is organized as follows. Section 2 introduces a representative non-private recommendation algorithm and some background knowledge. Section 3 presents the detailed design of the hybrid approach. Section 4 discusses the experimental results and Sect. 5 concludes the paper.

2 Preliminaries

2.1 Recommendation Algorithm Without Privacy Guarantee

We first describe a recommendation algorithm [12] without privacy guarantee. Suppose there are n users and m items. Based on the data provided by n users, the recommender system has two matrices in hand. One is a rating matrix $R_{n \times m}$ that contains the ratings of n users for m items where r_{ui} indicates the rating of user u for item i . The other auxiliary (binary) matrix $E_{n \times m}$ indicates the presence of ratings, where $e_{ui} = 1$ means u has rated for i and $e_{ui} = 0$ means u does not. The two matrices are the input to the recommendation algorithm, while the output is predicted ratings of items that users have not rated.

Some users tend to give higher ratings than other users, and some items tend to receive higher ratings than others. This difference will make the recommendation result disappointing, so it is necessary to subtract user effects and item effects from ratings. We first compute the global average of $R_{n \times m}$:

$$GAvg = \frac{\sum_R r_{ui}}{\sum_E e_{ui}}$$

Then, we center ratings by computing and subtracting average ratings for items and users:

$$r'_{ui} = r_{ui} - UAvg(u)$$

$$UAvg(u) = \frac{\sum_i (r_{ui} - IAvg_i) + \beta_u \cdot GAvg}{\sum_i e_{ui} + \beta_u}, \quad IAvg_i = \frac{\sum_u r_{ui} + \beta_m \cdot GAvg}{\sum_u e_{ui} + \beta_m}$$

where $IAvg$ and $UAvg$ are dampened by β_m and β_u fictitious ratings of the global average, respectively. Here, β_m is the average number of ratings for item m , and β_u is the average number of rating items for user u .

Finally, we use the centered ratings to calculate the covariance matrix, which indicates the relationships between items:

$$Cov(i,j) = \frac{\sum_u w_u r'_{ui} r'_{uj}}{\sum_u w_u e_{ui} e_{uj}}$$

where w_u is per-user weights equaling to the reciprocal of $\|e_u\|$. The final recommendation result can be made by passing this covariance matrix to a large number of advanced learning and prediction algorithms, such as the k -nearest neighbor (kNN) method proposed by Bell and Koren [2].

2.2 Differential Privacy (DP)

Intuitively, differential privacy means the probability an attacker who is able to observe the computation’s output learns any record’s presence in or absence from the computation’s input should be indistinguishable [3,4]. The formal definition is as follows:

Definition 1. *A randomized function f provides ϵ -differential privacy if for any neighboring data bases A and B ($A \triangle B = 1$), and any subset S of possible outcomes $Range(f)$,*

$$Pr[f(A) \in S] \leq exp(\epsilon) \times Pr[f(B) \in S]$$

Two datasets A and B are adjacent if there is only one individual record difference between them ($A \triangle B = 1$). The parameter ϵ is the privacy budget, which can be used to control the level of privacy protection. The smaller the value of ϵ is, the stronger privacy protection it provides. DP guarantees the output is insensitive to any individual record. The probability that an attacker can correctly guess whether or not an individual record is in the dataset is at most $exp(\epsilon)$ based on the outputs of calculations. It satisfies a *composability property* defined as follows: The sequence of $f_i(A)$ provides $(\sum_i \epsilon_i)$ -differential privacy, where f_i each provides ϵ_i -differential privacy. Therefore, the ϵ parameter can be considered as an accumulative privacy cost as more steps are executed. These costs keep accumulating until they reach an allotted privacy budget.

A common way to obtain differential privacy is by applying random noise to the measurement. The amount of noise added depends on the L_1 -sensitivity of the evaluated function, which is the largest possible change in the measurement given a change in a single record in the dataset. In general, the L_k -sensitivity of a function f is given by:

$$S_k(f) = \max_{(A \triangle B = 1)} \|f(A) - f(B)\|_k$$

where $\|\cdot\|_k$ denotes the L_k -norm.

Given a function $f:D \rightarrow \mathbb{R}^d$, Laplace mechanism obtains ϵ -differential privacy by adding noise sampled from Laplace distribution, with a calibrated scale $b = S_1(f)/\epsilon$. The following computation maintains ϵ -differential privacy:

$$K(x) = f(x) + (Laplace(S_1(f)/\epsilon))^d$$

2.3 Randomized Perturbation

The basic idea of randomized perturbation is to perturb the data in such a way that certain computations can be done while preserving users’ privacy. Although data from each user is scrambled, if the number of users is significant large, the aggregate information of these users can be estimated with decent accuracy. Such property is very useful for computations that are based on aggregate information. Scaler product and sum are among such computations.

Let r^a and r^b be the original vectors, where $r^a = (r_1^a, \dots, r_i^a)$ and $r^b = (r_1^b, \dots, r_i^b)$. r^a is disguised by $v^a = (v_1^a, \dots, v_i^a)$, and r^b by $v^b = (v_1^b, \dots, v_i^b)$, where v^a and v^b are uniformly distributed in domain $[-\gamma, \gamma]$. Let $r'^a = r^a + v^a$ and $r'^b = r^b + v^b$ be disguised data that are known. Because v^a and v^b are uniformly distributed, the scalar product of r^a and r^b can be estimated from r'^a and r'^b and the sum of the values of r^a can be estimated from r'^a as follows:

$$\sum_{i=1}^n (r_i + v_i) = \sum_{i=1}^n r_i + \sum_{i=1}^n v_i \approx \sum_{i=1}^n r_i \tag{1}$$

$$r'^a \cdot r'^b = \sum_{i=1}^n (r_i^a r_i^b + r_i^b v_i^a + r_i^a v_i^b + v_i^a v_i^b) \approx \sum_{i=1}^n r_i^a r_i^b \tag{2}$$

3 The Hybrid Approach

Figure 1 shows the whole life-cycle of a typical recommender system armed with our hybrid privacy-preserving approach. Three stages are involved in the process of recommendation: *data collection*, *data publication* and *data prediction*. In the first stage, users’ original data are disguised through randomized perturbation, resulting in perturbed rating matrix R and auxiliary matrix E . Based on the two perturbed matrices, the recommender system computes global average, item averages, user averages, and finally the covariance matrix for data publication. All these data are masked with particular amount of noise to guarantee differential privacy. With the added noise, the covariance matrix is ready for publication and can be fed into an existing learning and prediction algorithm (e.g., the k NN method [2]) with no changes. As mentioned earlier, the challenge here is how to ensure recommendation accuracy and realize differential privacy on the perturbed data effectively.

3.1 Methodology

In the data collection stage, the recommender system decides on a range $[-\gamma, \gamma]$ and let each user know. Then, each user u disguises her ratings r_{ui} by adding noise that is uniformly distributed in the domain $[-\gamma, \gamma]$. The recommender system collects these disguised data r'_{ui} to form two perturbed matrices.




Fig. 1. Overview of the hybrid approach for privacy-preserving recommender systems

In the stage of data publication, the recommender system injects noise to three different values: the global average, the per-item average and the covariance matrix. Note that, the noises for these values are different and depend on the sensitivity of the underlying functions of computing these values. For global average, Laplace distributed noise is added to guarantee its privacy:

$$GAvg = \frac{\sum_R r'_{ui} + Laplace(\Delta r_1/\epsilon_1)}{\sum_E e_{ui} + Laplace(\Delta r'_1/\epsilon_1)}$$

where Δr_1 and $\Delta r'_1$ are the sensitivity of function $\sum_R r'_{ui}$ and $\sum_E e_{ui}$, respectively. Their exact values will be discussed in the next subsection. We then use the global average to produce a stabilized per-item average rating by β_m at value $GAvg$ for each item:

$$IAvg_i = \frac{(\sum_u r'_{ui} + Laplace(\Delta r_2/\epsilon_2)) + \beta_m \cdot GAvg}{(\sum_u e_{ui} + Laplace(\Delta r'_2/\epsilon_2)) + \beta_m}$$

where Δr_2 and $\Delta r'_2$ are the sensitivity of $\sum_u r'_{ui}$ and $\sum_u e_{ui}$, respectively.

Having published the average rating for each item, we center the ratings for each user as follows, taking shrinking parameter β_u at global average, where c_u is the number of ratings by user u :

$$UAv(u) = \frac{\sum_i (r'_{ui} - IAvg_i) + \beta_u \cdot GAvg}{c_u + \beta_u}$$

We subtract user effects average from the appropriate ratings and clamp the resulting centered ratings to the intervals $[-B, B]$, to lower the sensitivity of the measurements at the expense of the relatively few remaining large entries:

$$\hat{r}_{ui} = \begin{cases} -B & \text{if } r'_{ui} - UAv(u) < -B \\ r'_{ui} - UAv(u) & \text{if } -B \geq r'_{ui} - UAv(u) < B \\ B & \text{if } r'_{ui} - UAv(u) \geq B \end{cases}$$

The final measurement we make of the private data is the covariance of the perturbed and clamped user ratings vectors. To retain the difference between

users, we take the non-uniform averages by using per-user weights w_u which equals to the reciprocal of $\|e_u\|$. Then, the covariance will be published as:

$$Cov(i,j) = \frac{\sum_u w_u \hat{r}_{ui} \hat{r}_{uj} + Laplace(\Delta r_3 / \epsilon_3)}{\sum_u w_u e_{ui} e_{uj} + Laplace(\Delta r'_3 / \epsilon_3)}$$

where Δr_3 and $\Delta r'_3$ are the sensitivity of $\sum_u w_u \hat{r}_{ui} \hat{r}_{uj}$ and $\sum_u w_u e_{ui} e_{uj}$, respectively.

3.2 Theoretical Analysis

As mentioned earlier, different functions have different sensitivities, which determines the amount of noise needed for differential privacy. In this subsection, we analyze the sensitivities of different functions involved in the data publication. The sensitivity values Δr_1 and Δr_2 are both $\tau + 2\gamma$, where τ is the maximum possible difference in raw ratings, and γ is the parameter of RP. For example, if the range of rating is from 1 to 5, the τ then equals to 4. From Theorem 2, the sensitivity value Δr_3 is $2B(\tau + 2\gamma) + 3B^2$. For $\Delta r'_1$ and $\Delta r'_2$, their values are both 1, because the maximum possible difference is 1 in the auxiliary matrix when e^a and e^b differ on only one value. The value of $\Delta r'_3$ is 3 which is clear from Theorem 3.

Theorem 1. *Let r^a and r^b differ on one rating, τ be the maximum possible difference in raw ratings. Considering the randomized perturbation before collecting the data, the maximum possible difference in the processed ratings is $\tau + 2\gamma$. For centered and clamped ratings \hat{r}^a and \hat{r}^b , we have*

$$\|\hat{r}^a - \hat{r}^b\|_1 \leq \tau + 2\gamma + B$$

Proof: If r^a and r^b are two sets of ratings which differ on one rating, present in r^b at r^b_{ui} , others are everywhere equal, except for the ratings of user u . For the ratings in common between r^a and r^b , the difference is at most the difference in the subtracted averages:

$$|U\text{Avg}(u)^b - U\text{Avg}(u)^a| = \frac{|r_{ui} - U\text{Avg}(u)^a|}{c_u^b + \beta_p} \leq \frac{\tau + 2\gamma}{c_u^b + \beta_p}$$

For the new rating r_{ui} , its previous contribution of zero is replaced with the new centered and clamped rating, at most B in magnitude. Hence, we have

$$\|\hat{r}^a - \hat{r}^b\|_1 \leq c_u^a \times \frac{\tau + 2\gamma}{c_u^b + \beta_p} + B$$

Note that $c_u^b = c_u^a + 1$ and the maximal value of c_u^a is $\beta_p + 1$. Therefore, the upper bound of $\|\hat{r}^a - \hat{r}^b\|_1$ is $\tau + 2\gamma + B$. \square

Theorem 2. Let r^a and r^b differ on one rating. Taking $w_u = 1/\|e_u\|_1$, we have

$$\|w_u^a r_{ui}^a r_{uj}^a - w_u^b r_{ui}^b r_{uj}^b\| \leq 2B(\tau + 2\gamma) + 3B^2$$

Proof: For the difference $w_u^a r_{ui}^a r_{uj}^a - w_u^b r_{ui}^b r_{uj}^b$, we can rewrite it as $w_u^a r_{ui}^a (r_{uj}^a - r_{uj}^b) + w_u^b (r_{ui}^a - r_{ui}^b) r_{uj}^b + (w_u^a - w_u^b) r_{ui}^a r_{uj}^b$, as $\|e_u^b - \|e_u^a\| \leq 1$, we have that

$$w_u^a - w_u^b = \frac{1}{\|e_u^a\|} - \frac{1}{\|e_u^b\|} \leq \frac{1}{\|e_u^a\| \|e_u^b\|}$$

The original matrix difference is bounded by

$$\left(\frac{\|r_i^a\|}{\|e_i^a\|} + \frac{\|r_i^b\|}{\|e_i^b\|} \right) \|r_i^a - r_i^b\| + \frac{\|r_i^a\| \|r_i^b\|}{\|e_i^a\| \|e_i^b\|}$$

Giving Theorem 2, we have the upper bound $2B(\tau + 2\gamma) + 3B^2$. □

Theorem 3. Let e^a and e^b differ on one rating presence or absence. Taking $w_u = 1/\|e_u\|_1$, we have

$$\|w_u^a e_{ui}^a e_{uj}^a - w_u^b e_{ui}^b e_{uj}^b\| \leq 3$$

Proof: Between the two weight matrices, similarly, we can rewrite it as $w_u^a e_{ui}^a (e_{uj}^a - e_{uj}^b) + w_u^b (e_{ui}^a - e_{ui}^b) e_{uj}^b + (w_u^a - w_u^b) e_{ui}^a e_{uj}^b$. Then, we have the bound as follows:

$$\|w_u^a e_{ui}^a e_{uj}^a - w_u^b e_{ui}^b e_{uj}^b\| \leq \left(\frac{\|e_i^a\|}{\|e_i^a\|} + \frac{\|e_i^b\|}{\|e_i^b\|} \right) \|e_i^a - e_i^b\| + \frac{\|e_i^a\| \|e_i^b\|}{\|e_i^a\| \|e_i^b\|} = 3$$

□

The above theorems show that the hybrid approach takes into account the effect of noise introduced by RP on the noise injected by DP. If we directly use DP without considering the noise of RP, we will obtain weaker privacy protection. This result is guaranteed by the following theorem.

Theorem 4. The hybrid approach can provide stronger privacy protection than DP when raw rating data are disguised by RP.

Proof: First note that the sensitivity of $\sum_R r'_{ui}$ and $\sum_R r_{ui}$ are $\tau + 2\gamma$ and τ , respectively. To provide ϵ_1 -differential privacy for the global average, the hybrid approach injects noise v based on $Laplace(\frac{\tau+2\gamma}{\epsilon_1})$. As DP does not consider the noise introduced by RP, it will inject noise v' based on $Laplace(\frac{\tau}{\epsilon_1})$. The noise v' on the disguised raw data can actually provide ϵ'_1 -differential privacy where $\frac{\tau+2\gamma}{\epsilon'_1} = \frac{\tau}{\epsilon_1}$. Thus we have $\epsilon'_1 = \frac{\tau+2\gamma}{\tau} \epsilon_1 > \epsilon_1$. Likewise, we can have $\epsilon'_2 > \epsilon_2$ and $\epsilon'_3 > \epsilon_3$ based on the sensitivity values given in Theorem 2, where ϵ'_2 and ϵ'_3 are the actual privacy budget DP can provide for item average and covariance, respectively. According to the composability property of differential privacy, we have $\epsilon = \epsilon_1 + \epsilon_2 + \epsilon_3$ and $\epsilon' = \epsilon'_1 + \epsilon'_2 + \epsilon'_3$. Clearly, $\epsilon < \epsilon'$, which completes the proof. □

We now examine the effect of the noise added by RP in the hybrid approach on recommendation accuracy. As mentioned earlier, though raw ratings from each user are perturbed, the aggregate information of these ratings can be estimated with decent accuracy if the number of users is significant large. Suppose $GAvg$ is the global average of the perturbed raw data collected in the hybrid approach and $GAvg^*$ is the global average of the raw data collected in DP. Clearly, we have

$$GAvg = \frac{\sum_R r'_{ui} + Laplace(\Delta r_1/\epsilon_1)}{\sum_E e_{ui} + Laplace(\Delta r'_1/\epsilon_1)}, \quad GAvg^* = \frac{\sum_R r_{ui} + Laplace(\Delta r_1^*/\epsilon_1)}{\sum_E e_{ui} + Laplace(\Delta r'_1/\epsilon_1)}$$

where Δr_1 , $\Delta r'_1$, and Δr_1^* are the sensitivity of function $\sum_R r'_{ui}$, $\sum_E e_{ui}$, and $\sum_R r_{ui}$, respectively. If R is sufficiently large, we have $\sum_R r'_{ui} \approx \sum_R r_{ui}$ as the noise injected into r_{ui} is uniformly sampled from $[-\gamma, \gamma]$. Besides, it is important to notice that $\sum_R r'_{ui} \gg Laplace(\Delta r_1/\epsilon_1)$. Thus, we have: $GAvg \approx GAvg^*$. Likewise, we can conclude that $Cov(ij) \approx Cov^*(ij)$, which indicates that the noise added by RP in the hybrid approach has limited effect on recommendation accuracy.

4 Experiments

4.1 Experimental Setting

In this section, we evaluate our hybrid approach for privacy-preserving recommender systems. As discussed earlier, both RP and DP introduce noise into recommendation computation, so it is worth studying the prediction accuracy when combining the two techniques. Therefore, we examined each of the three methods (i.e., RP, DP, and the hybrid one) in turn to see its effect on recommendation accuracy. All experiments were conducted on three real world datasets: Netflix³ consists of roughly 100 M ratings of 17770 movies contributed by 480 K users; MovieLens⁴ consists of 100 K ratings of 1682 movies contributed by 943 users; Yahoo⁵ consists of 23 M ratings of 11915 movies contributed by 7742 users. The rating of the three datasets are all from 1 to 5.

By adjusting the parameters of the noise distributions we use (i.e., γ of RP and ϵ of DP), our approach provides different randomized perturbation and differential privacy guarantees, and consequently, the recommendation outputs have different accuracy values. In our experiment, the recommendation accuracy is measured by the *root mean squared error* (RMSE) on the test datasets:

$RMSE = \sqrt{\frac{\sum_X (x-x')^2}{|X|}}$ where X consists of all values needs to be predicted in the test set and $|X|$ is the size of X , x' is the predicted value and x is the original value in the test set. A smaller RMSE value indicates a more accurate recommendation result. Regarding the training set and test set, the MovieLens

³ <http://www.netflixprize.com>.

⁴ <http://grouplens.org/datasets/movielens>.

⁵ <https://webscope.sandbox.yahoo.com>.

data is divided into two parts, 80% for the training set and 20% for the test set. For Netflix data, the test set is the Probe set. For Yahoo dataset, the training set contains 7642 users and the test set has 2309 users. The test set is gathered chronologically after the training set.

We applied the k NN method [2] to the covariance matrix for the final recommendation. The value of k is fixed at 20, and the clamping parameter B is set to 1. Following the work in [12], for any ϵ , we set the respective ϵ_i as follows: $\epsilon_1 = 0.02 \times \epsilon, \epsilon_2 = 0.19 \times \epsilon, \epsilon_3 = 0.79 \times \epsilon$. All experiments were conducted on a Dell PowerEdge R930 server which is equipped with 2.2 GHz CPU and 2 TB RAM. Each experiment was run 10 times and the average results were reported.

4.2 Experimental Results

4.2.1 RP’s Effect on Accuracy. Figure 2 shows the recommendation accuracy when γ increases from 0.5 to 3.5 with a step of 0.5. Clearly, the accuracy decreases on all three datasets. This is because the noise added by RP is determined by the parameter γ . In particular, the larger the γ is, the wider range the random noise is in. Therefore, more randomness is likely to be added into the original data, resulting in less accurate recommendation.

4.2.2 DP’s Effect on Accuracy. Figure 3 shows the recommendation accuracy when ϵ increases from 0.1 to 10. From the results, we can see that the accuracy decreases rapidly when DP provides strong privacy guarantee (i.e.,

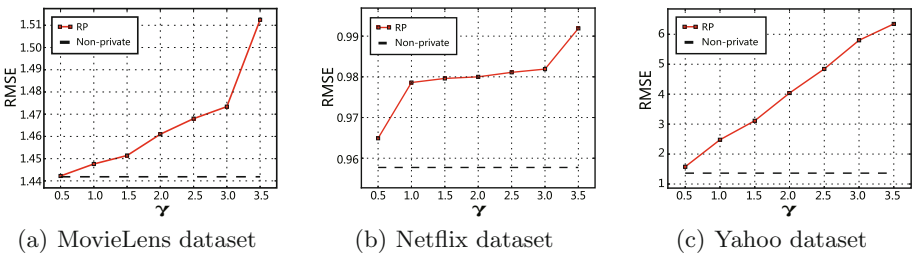


Fig. 2. RP’s effect on recommendation accuracy




Fig. 3. DP’s effect on recommendation accuracy

when $\epsilon < 1$). When providing a weak privacy protection (i.e., when $\epsilon > 1$), the accuracy approaches to a constant. Such observations imply that a large ϵ contributes little to the accuracy but weakens the privacy protection.

4.2.3 Effect of the Hybrid Approach on Accuracy. We first examine the hybrid approach by taking DP as the baseline. Figure 4 depicts how the recommendation accuracy of DP is affected by RP. It is clear that no matter which γ is used in RP, the overall trend of DP remains the same, that is, the accuracy decreases as ϵ approaches to 0, indicating a stronger privacy guarantee. Besides, when DP and RP work together, larger γ often leads to less accuracy. For example, DP plus RP with $\gamma = 0.5$ is more accurate than DP plus RP with $\gamma = 3.5$ on MovieLens dataset. Finally, it is worth noting that in most cases the combination of DP and RP makes recommendation less accurate, which coincides with our common sense as both of them introduce noise into the original data. However, their combination sometimes results in a win-win situation where both the accuracy and the privacy becomes better, as seen in Fig. 4(b). To make this clear, we draw in Fig. 5 the RMSE ratio between the hybrid approach and DP. We can see that for Netflix dataset, the combination of DP and RP sometimes outperforms DP only, especially when γ is small, for example, 0.5. Besides, the accuracy loss of DP plus RP is acceptable on MovieLens and Netflix datasets, but is not satisfactory on Yahoo dataset. A possible reason might be that MoveiLens and Netflix datasets have similar rating distribution, which is different from Yahoo dataset. We then examine the hybrid approach by taking RP as the baseline.




Fig. 4. RP’s effect on DP




Fig. 5. RMSE ratio between the hybrid approach and DP




Fig. 6. DP's effect on RP




Fig. 7. RMSE ratio between the hybrid approach and RP

Figure 6 depicts how the recommendation accuracy of RP is affected by DP. We can again see that the overall trend of RP remains the same (i.e. the accuracy decreases as γ increases) no matter which ϵ is used in DP. Besides, smaller ϵ is more likely to make the recommendation less accurate, as stronger privacy guarantee is provided in DP through injecting more noise into the data. We also notice that, for MovieLens dataset, the combination of RP and DP makes recommendation less accurate, but in an acceptable range. For the Netflix dataset, however, their combination is indeed a good choice as we can obtain additional privacy guarantee while not sacrificing recommendation accuracy. Further note that this is true for any combination of γ and ϵ in our experiments, as shown in Fig. 7(b). The accuracy loss is still unsatisfactory on Yahoo dataset, especially when DP provides strong privacy guarantee, as depicted in Fig. 7(c).

4.2.4 Efficiency of the Hybrid Approach. Figure 8 shows the running time of the hybrid approach. It is clear that the running time increases when the rating matrix becomes large, but the total computation cost is acceptable even on a moderate server. In particular, for the MovieLens dataset where the size of rating matrix is about $1000 * 1700$, the hybrid approach only needs 33s. Even for large Netflix dataset whose rating matrix is $380 K * 500$, the hybrid approach can be completed within less than 2.5 h. The computation cost of the hybrid approach mainly comes from DP, as RP only requires few simple operations and is done at user side. Thus, the hybrid approach has the same computation complexity as DP, but can provide stronger privacy guarantee than DP as shown in Theorem 4.

4.2.5 The Hybrid Approach Vs DP. Figure 9 depicts the accuracy comparison of the hybrid approach and DP when the raw rating data are disguised by RP with different γ . The privacy budget ϵ is set to 1 in both methods. In all datasets, we can see that DP has a better performance than the hybrid approach. This, however, exactly shows DP cannot provide sufficient privacy guarantee over the perturbed data, as it underestimates the sensitivity of the functions on the perturbed data. This result coincides with Theorem 4, which says the hybrid approach can provide stronger privacy protection than DP over the data disguised by RP. Further, the RMSE difference of the two methods is small, which means the hybrid approach has an acceptable accuracy loss while providing stronger privacy guarantee.

4.2.6 Summary. From the above discussion, we can see that, by carefully injecting appropriate noise into the perturbed data based the sensitivity analysis of different functions involved in the recommendation computation, the hybrid approach can provide more privacy protection with acceptable accuracy loss. More interestingly, the hybrid approach will not necessarily lead to less recommendation accuracy, which initially contradicts our common sense but has been validated subsequently by experiments on Netflix dataset, which is the largest dataset in our experiments. Besides, the integration of DP and RP does not affect their original trend of the relation between accuracy and privacy, which is also appealing as we still have control of the balance between accuracy and privacy in the hybrid approach.




Fig. 8. Efficiency of the hybrid approach




Fig. 9. RMSE of the hybrid approach and DP over perturbed data

5 Conclusion

We have presented a hybrid approach for privacy-preserving recommender systems by combining randomized perturbation (RP) and differential privacy (DP), which is more privacy-friendly than existing works as the user's private data are protected by randomized perturbation and no one can infer any single user's data from the normal recommendation output thanks to differential privacy. We have theoretically shown the noise added by RP has limited effect on recommendation accuracy and the noise added by DP can be well controlled based on the sensitivity analysis of functions on the perturbed data. We have conducted extensive experiments on real datasets and concluded that the combination of DP and RP is feasible not only in theory but also in practice.

Acknowledgment. This work was done while the first author was a visiting student at King Abdullah University of Science and Technology (KAUST). Research reported in this publication was partially supported by KAUST and Natural Science Foundation of China (Grant Nos. 61572336, 61572335, 61632016, 61402313).

References

1. Badsha, S., Yi, X., Khalil, I.: A practical privacy-preserving recommender system. *Data Sci. Eng.* **1**(3), 161–177 (2016)
2. Bell, R.M., Koren, Y.: Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In: *ICDM*, pp. 43–52 (2007)
3. Dwork, C.: Differential privacy: a survey of results. In: *International Conference on Theory and Applications of Models of Computation*, pp. 1–19 (2008)
4. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) *TCC 2006*. LNCS, vol. 3876, pp. 265–284. Springer, Heidelberg (2006). doi:[10.1007/11681878_14](https://doi.org/10.1007/11681878_14)
5. Elgamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory* **31**(4), 469–472 (1985)
6. Erkin, Z., Veugen, T., Toft, T., Legendijk, R.L.: Generating private recommendations efficiently using homomorphic encryption and data packing. *IEEE Trans. Inf. Forensics Secur.* **7**(3), 1053–1066 (2012)
7. Guerraoui, R., Kermarrec, A.M., Patra, R., Taziki, M.: D 2 p: distance-based differential privacy in recommenders. *VLDB* **8**(8), 862–873 (2015)
8. Huang, Y., Evans, D., Katz, J., Malka, L.: Faster secure two-party computation using garbled circuits. In: *USENIX Security Symposium*, vol. 201 (2011)
9. Liu, S., Liu, A., Liu, G., Li, Z., Xu, J., Zhao, P., Zhao, L.: A secure and efficient framework for privacy preserving social recommendation. In: Cheng, R., Cui, B., Zhang, Z., Cai, R., Xu, J. (eds.) *APWeb 2015*. LNCS, vol. 9313, pp. 781–792. Springer, Cham (2015). doi:[10.1007/978-3-319-25255-1_64](https://doi.org/10.1007/978-3-319-25255-1_64)
10. Ma, X., Li, H., Ma, J., Jiang, Q., Gao, S., Xi, N., Lu, D.: Applet: a privacy-preserving framework for location-aware recommender system. *Sci. China Inf. Sci.* **60**(9), 092101 (2017)
11. Machanavajjhala, A., Korolova, A., Sarma, A.D.: Personalized social recommendations: accurate or private. *VLDB* **4**(7), 440–450 (2011)

12. McSherry, F., Mironov, I.: Differentially private recommender systems: building privacy into the netflix prize contenders. In: KDD, pp. 627–636 (2009)
13. Nikolaenko, V., Ioannidis, S., Weinsberg, U., Joye, M., Taft, N., Boneh, D.: Privacy-preserving matrix factorization. In: CCS, pp. 801–812 (2013)
14. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999). doi:[10.1007/3-540-48910-X_16](https://doi.org/10.1007/3-540-48910-X_16)
15. Polat, H., Du, W.: Privacy-preserving collaborative filtering using randomized perturbation techniques. In: ICDM, pp. 625–628 (2003)
16. Yao, A.C.C.: How to generate and exchange secrets. In: FOCS, pp. 162–167 (1986)
17. Zhan, J., Hsieh, C.L., Wang, I.C., Hsu, T.S., Liau, C.J., Wang, D.W.: Privacy-preserving collaborative recommender systems. *IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.)* **40**(4), 472–476 (2010)
18. Zhang, S., Ford, J., Makedon, F.: Deriving private information from randomly perturbed ratings. In: SDM, pp. 59–69 (2006)