Local Differentially Private Matrix Factorization For Recommendations

1st Neera Jeyamohan, 2nd Xiaomin Chen, 3rd Nauman Aslam
Dept. of Computer and Information Sciences
Northumbria University
Newcastle, UK
jeyamohan.neera@northumbria.ac.uk

***Abstract*—In recent years recommendation systems have become popular in the e-commerce industry as they can be used to provide a personalized experience to users. However, performing analytics on users' private information has also raised privacy concerns. Therefore, various privacy protection mechanisms have been proposed for recommendation systems. Yet most of these methods provide privacy protection against user-side adversaries and disregards the privacy violations caused by the service providers. In this paper, we propose a local differential privacy mechanism for matrix factorization based recommendation systems. In the proposed method, users perturb their ratings locally on their devices using Laplace and randomized response mechanisms and send the perturbed ratings to the service provider. We evaluate the proposed mechanism using Movielens dataset and demonstrate that it can achieve a satisfactory trade-off between data utility and user privacy.**

***Index Terms*—Local Differential Privacy, Matrix Factorization, Recommendation System, Laplace Mechanism, Randomized Response**

I. Introduction

Increased adoption of mobile devices has produced extensive transformation in the e-commerce industry. As a result, users are more inclined to use online shopping platforms compared to traditional in-person shopping. Many e-commerce platforms use recommendation systems to aid users to find new items that are quite likely of interest to them. Recommendation systems are proven to support businesses to enhance user satisfaction, build customer loyalty and even form new user habits [1]. However, since recommendation systems are aggregating and analyzing data related to user's activities over e-commerce platforms, they often raise privacy concerns.

Differential Privacy is a relatively new privacy model which guarantees that the observation of an output of a randomized mechanism should not allow an adversary to infer about an instance in the dataset [2]. Compared to conventional privacy-preserving methods, e.g. cryptography, anonymisation, it has been proven that differential privacy based methods can provide strong privacy protection in cases where the adversary possesses a substantial amount of auxiliary information about users [2]. Differential privacy based collaborative filtering recommendation mechanisms have been initially proposed in [10], [11]. However, almost all of the existing differential privacy based methods consider privacy attacks launched by third-party adversaries. They assume that the e-commerce Service Provider (SP) is trusted and there is no risk of data misuse by them. Unfortunately many SPs are inclined to collect more data than required, which raises privacy concerns over data aggregation and analytics carried out by them. [5] shows the risk of persistent data aggregation by SPs and how analyzing an individual's historical ratings can reveal sensitive information such as user's political preference, medical conditions and even religious disposition.

In this work, we propose a local differentially private matrix factorization based recommendation mechanism. Instead of perturbing the aggregated ratings or output of any query related to recommendation system, our mechanism perturbs the original ratings of the users locally in their devices before forwarding it to the SP. Hence, the SP will be able to aggregate only the perturbed ratings. The randomness in the perturbed ratings will cause more prediction errors in the recommendation. Therefore, we use randomized response mechanism which is concatenated with Laplace mechanism locally at user-side to reduce the degree of perturbation without sacrificing privacy. The ratings aggregated at the SP are still perturbed data, but the randomized response mechanism provides the freedom to tune the differential privacy budget in terms of the recommendation accuracy requirement. This approach improves accuracy with the plausible guarantee of deniability of actual rating. The aim of this work is to guarantee strong privacy protection for each user and at the same time assure significant recommendation accuracy for service providers. We demonstrate that our mechanism can achieve a suitable trade-off between the recommendation accuracy and the strength of privacy protection. The user will be able to tune the local differential privacy budget in terms of their requirements.

This paper is organized as follows. Section 2 provides background information on differential privacy based recommendation mechanisms. In section 3, we briefly introduce local differential privacy, global sensitivity and Laplace mechanism as preliminaries. In section 4 , we explain our proposed local differentially private recommendation system. In Section 5, we demonstrate the performance of our proposed model with experiments carried out on Movielens dataset. Finally, we provide conclusions that are drawn from our work.

978-1-7281-2741-5/19/$31.00 ©2019 IEEE

II. Related Work

Cryptography and annonymisation based techniques were first used to protect privacy in recommendation systems. Cryptography based techniques are applicable and useful only in distributed recommendation systems [3]. However, these mechanisms rely heavily on complex computations, making them non-desirable to be used in recommendation systems. Anonymisation based techniques [4] are often used to hide the relationship between a user and his profile by removing any personally identifiable information. Although it does not require a complex computation, it cannot protect users from adversaries who possess substantial background knowledge about the users and the system. Both cryptography and anonymisation based solution can not defend against privacy attacks caused by service providers.

Since privacy concerns have been raised more frequently, differential privacy based models have been proposed for recommendation services through number of earlier works. The first privacy protection model for recommendation systems based on differential privacy [10] demonstrates that achieving higher data utility under strong privacy guarantee in neighbourhood based and model based collaborative filtering systems is possible. Following this work, another work [11] proved that the recommendation accuracy can be further increased by using an exponential mechanism. However, all these approaches are applicable under the assumption that the service provider is trustworthy.

Local differential privacy based solutions are introduced in many applications to prevent privacy violations caused by service providers. A number of works have adopted local differential privacy in recommendation systems. For instance, [6] proposes a mechanism where user's ratings are perturbed within pre-defined item category. Even though the proposed mechanism provides protection against an untrusted service provider, it can still reveal a user's preferred item category. Another proposed model [7] introduces noise through perturbing a single rating of the user rather than masking all their ratings through an objective function of matrix factorization. The work also suggested that this model will allow the user to either leave or join in the recommendation process without letting the service provider knowing about it. In our work, we apply local differential privacy in such manner user's every single rating is perturbed. Although this approach introduces a huge noise to aggregated data, we show that the recommendation accuracy can be still improved.

III. Preliminaries

A. Local Differential Privacy

Earlier work on differential privacy has established a strong privacy protection model for algorithms employed on aggregated databases assuming a trusted aggregator will curate private information of users. In practice, users are reluctant to share their information with a trusted aggregator. Local differential privacy based algorithms are introduced to capture the requirements of users whose private information can be released independently to a data aggregator without worrying about privacy violations. In Local differential privacy models each user is required to perturb their own data before releasing it to a data aggregator.

Definition 1: (Local Differential Privacy) A Non-deterministic randomized function M is ε-deferentially private for all inputs v_1 and v_2 and for all possible subsets of r of the output domain R if:

$$Pr[M(v_1) = r] \leqslant e^{\varepsilon} \times Pr[M(v_2) = r] \tag{1}$$

Definition 1 ensures that when the data aggregator aggregates all the outputs of a user, he would not be able to distinguish whether r is an outcome of input v_1 or v_2. The parameter ε is called privacy budget which is used as a privacy loss measure of the randomized mechanism and can be adjusted according to the privacy requirement. When ε increases then the privacy loss incurred by the mechanism also increases. In local differential privacy model each user can operate under separate privacy budget ε. But for simplicity we have assumed all users would be sharing the same privacy budget.

B. Global Sensitivity

The global sensitivity would be used to measure the maximum change caused in query outputs for any two neighboring datasets.

Definition 2: (Global Sensitivity) Global Sensitivity of any randomized function $f : X^n \rightarrow R$ on two adjacent datasets D_1, D_2 differing by at most one element is

$$\Delta f = max_{D_1,D_2} \|f(D_1) - f(D_2)\|_1 \tag{2}$$

C. Laplace Mechanism

Definition 3: (Laplace Mechanism) For a given randomized function f and Dataset D such that $f : D \rightarrow R$, the randomized mechanism M will provide ε differential privacy if:

$$M(D) = f(D) + Lap(\Delta f/\varepsilon) \tag{3}$$

in which Δf is the global sensitivity of the randomized function and ε is the corresponding privacy budget.Laplace mechanism introduced in [2], adds appropriate noise to a query output. As the name suggested, the Laplace mechanism adds noise that is sampled from Laplace distribution.

IV. System Model

In this work, the service provider is considered untrusted and the users are not willing to share their ratings with them. The proposed local perturbation mechanism perturbs the user's original ratings locally before sending it to the service provider for aggregation. Then the perturbed ratings are used to train the matrix factorization prediction model. In our work, single value decomposition (SVD) algorithm [12] is used as the matrix factorization prediction model. Fig. 1 shows the proposed local differential privacy based collaborative filtering recommendation system. In Table I, we list the notations we used throughout this section.

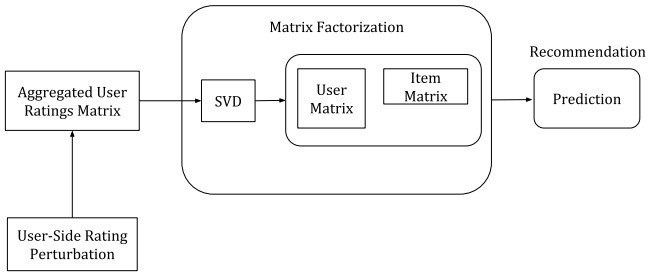

Fig. 1. Local differential privacy based collaborative filtering recommendation system

A. Local Differential Privacy Perturbation Mechanism

Our proposed system requires a local perturbation mechanism that runs on the user side. The perturbation mechanism will be a combination of Laplace and randomized response mechanisms. Since recommendation systems tend to use different scales to indicate the opinions of users on items, we adopt a Min-Max scaling approach to normalize the ratings. The normalized ratings are perturbed by adding Laplace noise derived from Laplace distribution. The range of ratings determines the global sensitivity of the Laplace mechanism. As the normalized ratings are in the range of $[0, 1]$, the global sensitivity should be $\Delta r = r_{max} - r_{min} = 1$. Let $R_{u,i}$ is the true rating given by user u for item i and ε_1 is the corresponding privacy budget for Laplace mechanism. Then the perturbed output of Laplace mechanism R_{LP} would be:

$$R_{LP} = R_{u,i} + Lap\left(\frac{1}{\varepsilon_1}\right) \tag{4}$$

However, as Laplace mechanism would introduce substantial randomness to original ratings, a randomized response mechanism will then be used to reduce the level of perturbation. Randomized response mechanism will report the true rating to the service provider with a predefined probability p. Let ε_2 be the privacy budget for the randomized response mechanism, then to guarantee that the randomized response mechanism satisfies ε_2 differential privacy the probability p would be defined as [2]:

$$p = \frac{e^{\varepsilon_2}}{1 + e^{\varepsilon_2}} \tag{5}$$

Sequential composability property of differential privacy makes it possible to compute the results of k differentially private mechanisms in sequence without compromising privacy while privacy budget ε is summed up at each step [2]. Using this property it can be proven that our proposed local rating perturbation mechanism is $(\varepsilon_1+\varepsilon_2)$-differentially private where ε_1 is the privacy budget for Laplace mechanism and ε_2 is the privacy budget for Randomized response mechanism.

B. Rating Prediction Model

Matrix factorization based recommendation methods are preferred among various collaborative filtering methods due to their prediction accuracy and computational scalability [8]. They allow the recommendation systems to identify latent features which are hidden behind the interactions between users and items. The intuition behind using matrix factorization in recommendation systems is that the ratings given by a user for an item is influenced by some latent features. When trying to identify these latent features, matrix factorization algorithm assumes that they number of latent features are lesser than the number of users and items.

TABLE I
NOTATIONS

Notation	Meaning
u	User u
i	Item i
R_{LP}	Perturbed output of Laplace Mechanism
$R_{u,i}$	User u's rating on item i
Δr	Global Sensitivity of rating perturbaiton
r_{max}	Maximum possible rating
r_{min}	Minimum possible rating
ε_1	Privacy budget of Laplace Mechanism
ε_2	Privacy budget of Randomized response mechanism
p	Probability of reporting true rating
$R_{x\times y}$	Rating matrix
$r_{u,i}$	User u's rating on item i
$U_{x\times k}$	User feature matrix
$I_{k\times y}$	Item feature matrix
U_x	User feature vector
I_y	Item feature vector
$\hat{r_{u,i}}$	estimated product of feature vectors
λ	Constant to regularize feature
e_{ui}	Minimized square error for each rating
γ	Rate of minimizing error

The input for a matrix factorization based collaborative filtering system would be a rating matrix $R_{X\times Y}$ which contains all the ratings of X number of users on Y number of items. In the rating matrix an element $r_{u,i}$ represents a rating of user u on item i. Let's assume there are k number of latent features that the recommendation system would like to discover. The matrix factorization algorithm will factorize the rating matrix into two matrices : The user feature matrix $U_{X\times K}$ and item feature matrix $I_{K\times Y}$. Each row in the user feature matrix U_X will represent the relationship between each user u and latent features. Likewise, each row in the item feature matrix I_Y will represent the relationship between each item i and latent features.

The factorization is done in such way the rating matrix $R_{X\times Y}$ can be represented as the product of user feature matrix $U_{X\times K}$ and item feature matrix $I_{K\times Y}$. Each rating $r_{u,i}$ in rating matrix can be estimated as shown below:

$$\hat{r_{u,i}} = U_X \cdot I_Y^T \tag{6}$$

To find the two latent feature matrices $U_{X\times K}$ and $I_{K\times Y}$ from ratings matrix $R_{X\times Y}$, Stochastic Gradient Descent (SGD) method is used [12]. In summary this method initializes the two latent feature matrices with some arbitrary values. Then it calculates the product of latent matrices and estimates the difference between the product value $\hat{r_{u,i}}$ and the real rating $r_{u,i}$. Then SGD tries to find the local minimum of the difference iteratively. In SGD, optimal features are learned by

evaluating the minimized squared error for each rating in the rating matrix as shown below:

$$e_{ui}^2 = (r_{ui} - \hat{r_{u,i}})^2 \tag{7}$$

It is essential to understand the direction in which the values has to be modified in each iteration to minimize the error. The gradient can be formulated as shown below:

$$U_X = U_X + \gamma e_{ui} I_Y$$

$$I_Y = I_Y + \gamma e_{ui} U_X \tag{8}$$

Here, γ is a constant value which represents the rate of minimizing error. An extended matrix factorization algorithm has been introduced with a regularization constant to avoid model over-fitting. The modified minimized squared error can be computed as shown below [8].

$$Min(u,i) = \sum_{r_{u,i} \in R} \left[(r_{ui} - \hat{r_{u,i}})^2 + \lambda(\|U_X\|^2 + \|I_Y\|^2) \right] \tag{9}$$

The constant λ is introduced to regularize the features and to avoid model over-fitting.

V. Experimental Analysis

A. Experimental Setup

The local differential privacy based matrix factorization system was implemented in Python. Experiments are conducted on a commodity laptop with Intel i5 2.40 GHz CPU and 8 GB memory. Movielens [9] dataset is used for evaluation. The dataset consists of 100,000 ratings from 1000 users on 1700 movie items. Root Mean Square (RMSE) values are used to evaluate the prediction accuracy. RMSE can be defined as follows:

$$RMSE = \sqrt{\frac{\sum_{i=0}^{n-1} (r_i - r_i^{'})^2}{n}} \tag{10}$$

in which r_i is the true rating, $r_i^{'}$ is the predicted rating and n is the total number of ratings in the test set.

B. Experimental Setup

We compare 3 matrix factorization based collaboration filtering recommendation systems by employing different mechanisms. The following notations will be used to represent these schemes in the remainder of this section:

- Laplace-SVD : Differentially private collaborative filtering based on SVD with Laplace mechanism
- Laplace-RR-SVD : Differentially private collaborative filtering based on SVD with Laplace and randomized response mechanisms
- NonDP-SVD : Non differentially private collaborative filtering based on SVD

1) Impact of randomized response on accuracy: Fig 2 illustrates the effect of randomized response on accuracy. We compare three schemes i) SVD based collaborative filtering with Laplace mechanism(Laplace-SVD), ii) SVD based collaborative filtering with Laplace and the randomized response mechanism (Laplace-RR-SVD), and iii) SVD based collaborative filtering without any privacy protection (NonDP-SVD). The x-axis represents the privacy budget of Laplace mechanism ε_1. The RMSE of the NonDP-SVD scheme does not change when ε_1 increases. However, the RMSE value of the other two schemes decreases when ε_1 increases. Because when ε_1 increases, as a result, the level of noise added to the aggregated ratings through the Laplace-SVD and Laplace-RR-SVD schemes decreases. The RMSE of Laplace-RR-SVD scheme is substantially lower opposed to the Laplace-SVD scheme. This is because when the probability p increases in the Laplace-RR-SVD scheme more true ratings are reported to the service provider, which in return increases the prediction accuracy.

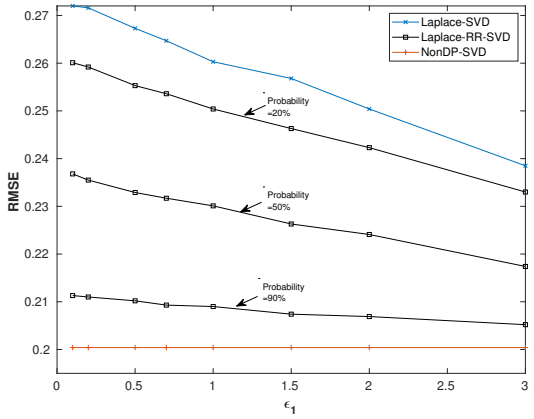

Fig. 2. RSME as ε_1 and ε_2 increase.

2) Privacy-Utility Tradeoff: Fig 3 plots the RMSE when ε_1 and ε_2 vary from 0.1 to 3 in step of 0.1. The result demonstrates the impact of the sequential combination of ε_1 and ε_2 on the prediction accuracy. As expected, when the privacy budgets ε_1 and ε_2 increase, the prediction accuracy also increases. It can be also seen that RMSE starts to converge when ε_1 and ε_2 approach to 3. This implies that a lower ε_1 and ε_2 values will restrict an adversary from learning about the existence of a user in a dataset. However, as ε_1 and ε_2 value lowers it also reduces the accuracy of recommendations. Nonetheless, when ε_1 and ε_2 values are higher, the proposed mechanism still provides guarantee the user the option of plausible deniability over an actual rating. We experimentally have shown that a proper trade-off decision can be made by choosing an appropriate ε_1 and ε_2 values, which in return would be beneficial for both users and service providers. Possible extension of our study should investigate on how to choose the upper and lower bounds of privacy budget ε through theoretical analysis.

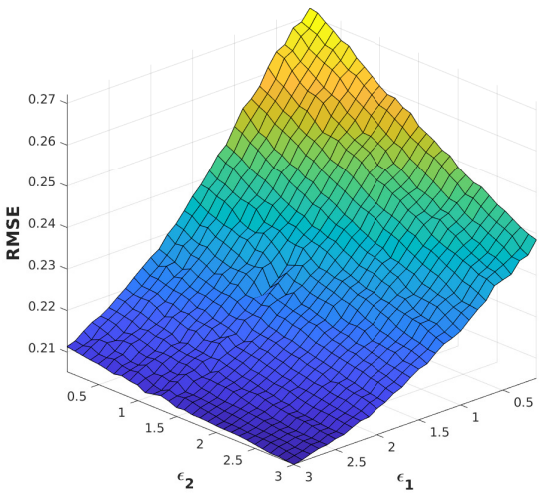

Fig. 3. RSME as ε_1 and ε_2 increase.

VI. Conclusion

This paper proposes a local differential private matrix factorization based collaborative filtering method which protects the privacy of user from an untrustworthy service provider. Local differential privacy is guaranteed by employing differentially private input perturbation mechanisms at the use's device. The original ratings of user is perturbed by Laplace and a randomized response mechanism. The evaluation of the proposed method demonstrates that the proposed input perturbation mechanism can achieve a satisfactory prediction accuracy while providing strong privacy protection. It also demonstrates that due to the introduction of randomized response mechanism, the accuracy loss caused by the Laplace mechanism can be effectively compensated. The proposed mechanism is designed to protect the privacy of user from service provider by perturbing their ratings. However, the mechanism currently does not hide the information about the items that are rated by the user. As in many cases the information about items purchases is a sensitive information as the ratings, the future work will address a mechanism that will protect both ratings and items that are rated.

References

[1] C. Xu, D. Peak and V. Prybutok, "A customer value, satisfaction, and loyalty perspective of mobile application recommendations", Decision Support Systems, vol. 79, pp. 171-183, 2015. Available: 10.1016/j.dss.2015.08.008.

[2] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy", Foundations and Trends in Theoretical Computer Science, vol. 9, no. 3-4, pp. 211-407, 2013. Available: 10.1561/0400000042.

[3] L. Sweeney, "k-ANONYMITY: A MODEL FOR PROTECTING PRIVACY", International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, vol. 10, no. 05, pp. 557-570, 2002. Available: 10.1142/s0218488502001648.

[4] C. John, "Collaborative filtering with privacy", in IEEE Symposium Security and Privacy, 2002, pp. 45-57.

[5] A.Narayanan and V.Shmatikov,"Robust de-anonymization of large datasets", in IEEE Symposium on Security and Privacy, 2008.

[6] Y. Shen and H. Jin, "Towards practical differentially private framework for personalized recommendation", in ACM SIGSAC conference on computer and communications security, 2016, pp. 180-191.

[7] J. Hua, C. Xia and S. Zhong, "Differentially private matrix factorization", in Proceedings of 7th International Conference in Artificial Intelligence, 2015, pp. 1763-1770.

[8] Y. Koren, R. Bell and C. Volinsky, "Matrix Factorization Techniques for Recommender Systems", Computer, vol. 42, no. 8, pp. 30-37, 2009. Available: 10.1109/mc.2009.263.

[9] F. Harper and J. Konstan, "The MovieLens Datasets", ACM Transactions on Interactive Intelligent Systems, vol. 5, no. 4, pp. 1-19, 2015. Available: 10.1145/2827872.

[10] F. Mcsherry and I. Mirnov, "Differentially private recommender systems: Building privacy into the netflix prize contenders", in Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, 2009, pp. 627-636.

[11] T. Zhu, G. Li, Y. Ren, W. Zhou and P. Xiong, "Differential privacy for neighborhood-based collaborative filtering.", in Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 2013, pp. 752-759.

[12] S. Badrul, K. George, K. Joseph and R. John, "Incremental singular value decomposition algorithms for highly scalable recommender systems", in Fifth international conference on computer and information science, vol. 27, 2002, pp. 28.