

A Local Differential Privacy based Hybrid Recommendation Model with BERT and Matrix Factorization

Jeyamohan Neera^a, Xiaomin Chen^b, Nauman Aslam, Biju Issac and Eve O'Brien

Department of Computer and Information Sciences, Northumbria University, U.K.

Keywords: Local Differential Privacy, Sentiment Analysis, Collaborative Filtering, Privacy.


Abstract: Many works have proposed integrating sentiment analysis with collaborative filtering algorithms to improve the accuracy of recommendation systems. As a result, service providers collect both reviews and ratings, which is increasingly causing privacy concerns among users. Several works have used the Local Differential Privacy (LDP) based input perturbation mechanism to address privacy concerns related to the aggregation of ratings. However, researchers have failed to address whether perturbing just ratings can protect the privacy of users when both reviews and ratings are collected. We answer this question in this paper by applying an LDP based perturbation mechanism in a recommendation system that integrates collaborative filtering with a sentiment analysis model. On the user-side, we use the Bounded Laplace mechanism (BLP) as the input rating perturbation method and Bidirectional Encoder Representations from Transformers (BERT) to tokenize the reviews. At the service provider's side, we use Matrix Factorization (MF) with Mixture of Gaussian (MoG) as our collaborative filtering algorithm and Convolutional Neural Network (CNN) as the sentiment classification model. We demonstrate that our proposed recommendation system model produces adequate recommendation accuracy under strong privacy protection using Amazon's review and rating datasets.


1 INTRODUCTION

Collaborative Filtering (CF) based recommendation systems predict a user's preference using similar users or relevant items. Even though CF algorithms produce satisfactory recommendation accuracy to an extent, they build upon the presumption that ratings reflect a user's true preferences and the actual quality of an item. This presumption does not always correspond with real scenarios, as impugned users might give low ratings and tolerant users tend to give high ratings. Such hypotheses play a huge role when a customer decides whether an item is suitable or unsuitable to fulfil their needs (Raghavan et al., 2012). Besides, users who give similar ratings to an item may have experienced distinct degrees of satisfaction (Cheng et al., 2018). Finally, CF algorithms often suffer from data sparseness problems due to a lack of ratings which results in low recommendation accuracy (Mobasher et al., 2007). Therefore, CF algorithms have started incorporating sentiment analysis to address these issues. Sentiment analysis is a technique

used to categorize text-based data to better understand users' attitudes and opinions in several domains. By combining sentiment analysis with a CF algorithm, the service provider can use ratings and reviews to further improve the recommendation accuracy. However, collecting both ratings and reviews imposes a higher risk of causing a violation of user privacy.

This paper proposes applying a Local Differential Privacy (LDP) based perturbation mechanism to a recommendation system that combines sentiment analysis with a CF algorithm to predict a user's preferences. We perturb the user's original ratings locally using a Bounded Laplace input perturbation mechanism (BLP) before sending them to the service provider. A deep learning-based sentiment analysis model is used to analyze user reviews. However, we propose that user reviews are tokenized locally and then sent to the service provider for aggregation and classification purposes. Since the service providers only aggregate the tokenized reviews, they cannot infer sensitive information about users without access to the original review data. Additionally, the perturbed ratings are used as the input sentiment labels, preventing the service provider from learning a user's actual sentiment using their tokenized review data. We use

^a  <https://orcid.org/0000-0001-8771-4193>

^b  <https://orcid.org/0000-0001-9267-355X>

Matrix Factorization (MF) with Mixture of Gaussian (MoG) as our CF algorithm. The MF-MoG model that runs at the service provider's end estimates the noise added to the aggregated perturbed ratings and predicts missing ratings simultaneously. The results of the empirical study show that our proposed recommendation model significantly improves the recommendation accuracy under a strong privacy guarantee.

2 RELATED WORK

Many works (Paterek, 2007; Pal et al., 2017) have investigated different approaches to improve the predictive performance of CF algorithms. Wang *et al.* (Wang et al., 2018) incorporated a CF recommendation system with a sentiment analysis model to obtain an optimised preliminary recommendation list first and then used it to produce a final recommendation list at the end. In another work, Osman *et al.* (Osman et al., 2019) proposed a recommendation system where the context of user's comments was taken into consideration and used to produce recommendations. Such approaches are more suitable when ratings are sparse and aid immensely in increasing recommendation accuracy.

Even though these proposed solutions improve the recommendation accuracy, they also cause serious privacy concerns in recommendation systems, and several works have proposed new ways to tackle these privacy concerns. Li and Sarkar (Li and Sarkar, 2011) presented an encryption-based solution to address the privacy concerns in a user-based CF recommendation system. McSherry and Mironov (McSherry and Mironov, 2009) proposed a differential privacy based solution in which noise is added to the item-to-item co-variance matrix. They proved that using differential privacy in recommendation systems offered strong privacy protection to users through experiments. Wang and Duan (Wang et al., 2016) proposed a privacy quantification model which synthesised an individual's influence and system privacy factors based on a user's perception. Similarly, Sutanto *et al.* (Sutanto et al., 2013) designed a personalised, privacy-safe application that enabled users to control their privacy. However, this method failed to produce accurate recommendations.

The aforementioned works consider the service providers to be trustworthy. Considering the existence of untrustworthy service providers, many works have begun to investigate the application of LDP in recommendation systems. LDP perturbs users' data on the user-side rather than on the service provider side. Liu *et al.* (Liu et al., 2015) proposed a recommendation

system where noise is added to users' ratings locally on the user-side through a randomised perturbation method. Meng *et al.* (Meng et al., 2018) in their work proposed using LDP based input perturbation mechanism only on ratings that were considered sensitive before sending them to the service provider. Shen and Jin (Shen and Jin, 2014) aimed to hide users' preference toward an item using an instance-based admissible mechanism on all the ratings. Shin *et al.* (Shin et al., 2018) proposed an LDP-based recommendation model which used a randomised response mechanism to add noise to users' ratings.

Most solutions that have been proposed in the literature to address privacy concerns in recommendation systems concentrate only on ratings based recommendation systems and have not yet addressed privacy concerns when user reviews are taken into consideration. Therefore, we propose using an LDP perturbation mechanism that perturbs users' ratings and tokenize reviews on the user-side to protect user's privacy from the service provider. We also propose estimating the noise added to the ratings at the server-side to improve the recommendation accuracy.

3 PRELIMINARIES

3.1 Local Differential Privacy

Each user perturbs their data locally before sending it to the service provider in the LDP setting. So the service provider collects only the perturbed data instead of the original data. This approach can significantly protect users' privacy from an untrustworthy service provider. Intuitively in LDP based settings, the data aggregator cannot infer whether a user's input x or x' produces the output y . Therefore, LDP offers plausible deniability to users.

Definition 1. A randomised mechanism M satisfies ϵ -LDP if for all possible pairs of user input x, x' and any subset y of all possible outcomes, we have the following inequality:

$$\Pr[M(x) \in y] \leq e^\epsilon \times \Pr[M(x') \in y].$$

The privacy budget ϵ acts as a metric of privacy loss at a perturbed data.

3.2 Bounded Laplace Mechanism

BLP perturbs data by ignoring any output values that do not fall inside a predefined domain and re-samples the noise from a Laplace distribution until the output value lies within the given bound. Unlike the Laplace

mechanism, BLP sanitises the perturbed output with bounding constraints (Holohan et al., 2018).

Definition 2. (*Bounded Laplace Mechanism*) Given a scale parameter b and a domain rating interval of (l, u) , the Bounded Laplace mechanism $M_{BLP} : R \rightarrow R^*$ is given by a conditional probability density function as follows:

$$f_W(r^*) = \begin{cases} \frac{1}{C_r(b)} \frac{1}{2b} e^{-\frac{|r^*-r|}{b}}, & \text{if } r^* \in [l, u], \\ 0, & \text{if } r^* \notin [l, u], \end{cases}$$

where $C_r(b)$ is a normalisation constant, r is the input to the BLP and r^* is the perturbed output.

The normalisation constant $C_r(b)$ is given as:

$$C_r(b) = 1 - \frac{1}{2} \left(\exp\left(-\frac{r-l}{b}\right) + \exp\left(-\frac{u-r}{b}\right) \right)$$

3.3 Bidirectional Encoder Representations from Transformers

The objective of sentiment analysis is to determine whether a user’s review communicates their positive or negative opinion. Deep learning-based techniques are proven to be highly effective in identifying user sentiments in several applications (Goularas and Kamis, 2019; Zarzour et al., 2021). Hence, we use BERT (Bidirectional Encoder Representations from Transformers) to create word embedding in our sentiment analysis model. BERT is a natural language processing model which was introduced by the Google AI team in 2018 (Devlin et al., 2018). BERT provides a contextualised representation, unlike other word-embedding models such as Word2Vec and GloVe, that generates a single representation for each word in a given input text. BERT uses transformer encoder layers to learn these contextual relations between words in a given text. The training of the BERT model takes place in two stages: pre-training and fine-tuning.

The BERT model is trained on an unlabelled corpus that contains text from English Wikipedia and Book Corpus dataset in the pre-training stage. These large collections of words allow BERT to capture extensive language knowledge. The resulting pre-trained model can then be fine-tuned for specific NLP tasks such as sentiment analysis. The fine-tuning stage is an essential step in training the BERT model. Even though pre-training produces a bidirectional unsupervised language representation of texts, fine-tuning allows this representation to be used in any NLP related tasks. In the fine-tuning stage, the BERT model is initialised with the same parameters as the pre-trained model, and the parameters are then adjusted according to the labelled data. BERT model

has several deployment types based on their model configurations, such as RoBERTa (Liu et al., 2019) which proposed a method to improve the training process of BERT and ALBERT (Lan et al., 2019) which reduces the model size through parameter sharing and factorising techniques.

4 SYSTEM DESIGN

In this section, we describe our proposed LDP based recommendation system that combines a CF algorithm with a sentiment analysis model and uses BLP as the input perturbation mechanism. The aim is to improve the recommendation accuracy while providing privacy protection to the users. Fig.1 illustrates the architecture of the proposed recommendation system. We assume that users can report their actual ratings and reviews anonymously so that the service provider can display these anonymous reviews and ratings on their platform. We do not discuss the architecture required for anonymous reporting as it is beyond the scope of this paper. We use the MF-MoG model as the CF algorithm that uses perturbed ratings as the input. The predicted rating combination module uses the outputs of the CNN classification model and the MF-MoG model to produce the final predicted rating.

Algorithm 1: BLP Mechanism for Noise Sampling.

- 1: **Input to the Mechanism: Original Rating**
 - 2: **Output of the Mechanism: Perturbed Rating**
 - 3: Generate noise from the Laplace distribution with mean 0 and variance of b
 - 4: Perturbed rating = Original rating + noise
 - 5: **If** Perturbed rating is in given domain interval:
 - 6: Perturbed rating is set
 - 7: **else**
 - 8: repeat Step 3
 - 9: **Return** Perturbed rating to SP
-

4.1 Input Rating Perturbation at User-side

First, users use BLP as an LDP perturbation mechanism to perturb their original ratings. The perturbed ratings are then sent to the service provider for aggregation. Algorithm 1 describes how a perturbed rating is generated using the BLP mechanism. It has been proven (Neera et al., 2021) that a sufficient condition for BLP to satisfy ϵ -local differential privacy in recommendation systems is when the local sensitivity is $\Delta f = l - u$ where l is minimum and u is the maximum rating in a given rating scale. The BLP mech-

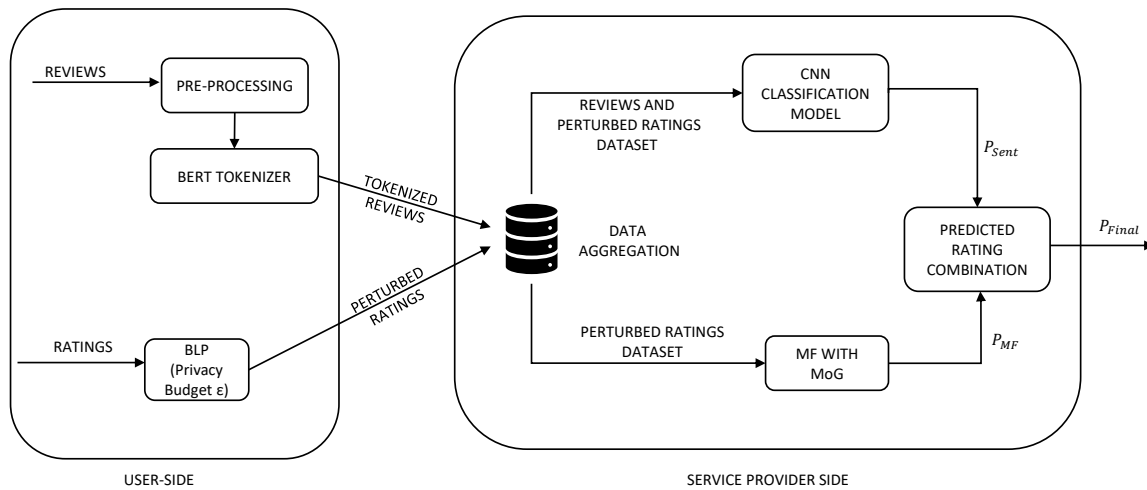


Figure 1: Proposed LDP-based Recommendation System with Sentiment Analysis and CF Algorithm.

anism ensures that the perturbed output rating is limited to the rating domain $[l, u]$ and still guarantees that the adversary cannot infer any information about the original rating by observing the perturbed rating.

4.2 Review Pre-processing at User-side

Sentiment analysis models require the input review data to be cleaned and processed before using them in a classification model. The original reviews of users are pre-processed and converted to a tokenized review before being sent to the service provider. First, words that lack relevant information, leading and trailing spaces, numbers, punctuation and stop words in the review text are removed. Additionally, the text is converted to lowercase. Then the cleaned review is split into individual words and then lemmatized. The lemmatization process converts the inflectional and derivational forms of a word to its common base form. For example, the words run, runs, and running are converted to the base word run. After lemmatization, we use BERT to compute the sequence embedding. BERT maps each word into a vector of numerical values so that words with similar meanings have a similar representation. Each user does the review pre-processing locally and sends only the numerical vector to the service provider so that the actual review text is never revealed to the service provider.

4.3 Multi Class Classification using CNN

In our work, we use BERT only as an encoder and a CNN model as the decoder to conduct sentiment classifications. Even though BERT itself can perform sentiment classification, the multi-label classification

layer has to be retrained on top of the transformer to perform sentiment prediction. Fig.2 illustrates the hybrid sentiment analysis model used in our recommendation system.




Figure 2: Hybrid sentiment analysis.

We use a CNN model for multi-class classification at the service provider’s end as it is proven to be effective in text review classification (Salinca, 2017). Since the CNN model uses perturbed ratings as the sentiment labels, we hide the true sentiment of the users from the service provider. This step adds another layer of privacy protection to the sentiment analysis model and offers plausible deniability to users. CNN model learns spatial hierarchies of features using three layers that is convolution, pooling and fully connected. The first two layers do feature extraction, and the final fully connected layer maps the extracted features into a relevant sentiment. We used a convolution layer with 256 filters and a linear rectification unit (ReLU) as the activation function. The output of the CNN model will be the predicted rating P_{Sent} .

4.4 Matrix Factorization with Mixture of Gaussian Model

We use Matrix factorization with a Mixture of Gaussian model (MF-MoG) (Neera et al., 2021) as the CF algorithm to make rating predictions on the ser-

vice providers' side. Since the CF algorithm uses perturbed ratings as the input, recommendation systems yield low recommendation accuracy. We use the Mixture of Gaussian (MoG) on the service provider side to estimate the noise added to the original ratings to enhance prediction accuracy. Since the post-processing property of LDP states that any further processing of a perturbed output of a differentially private mechanism does not violate the differential privacy principles (Dwork et al., 2014), the MF-MoG model still can provide privacy protection to users. Fig. 3 illustrates the MF-MoG recommendation system. The output of the MF-MoG model would be the predicted rating P_{MF} . Algorithm 2 details how the MF-MoG model estimates noise and predicts missing ratings.

Algorithm 2: Noise Estimation and Rating Prediction Model.

- 1: **Input:** Perturbed Ratings (R^*)
 - 2: **Output:** User latent factor U and Item latent factor V
 - 3: **Initialisation:** Model parameters U, V, Π and Σ are randomly initialised where Π is the Gaussian mixture proportion and Σ is the standard deviation.
 - 4: In E-step the posterior responsibility $\gamma_{ijk}^{(x)}$ is estimated
 - 5: **For** Until convergence
 - 6: (M-Step for updating $\Sigma^{(x+1)}$ and $\Pi^{(x+1)}$) Model parameters Σ and Π are computed.
 - 7: (M-Step for estimating V and U) Model parameters U and V are updated.
 - 8: (E-step for posterior responsibility γ_{ijk}) posterior responsibility γ_{ijk} is computed using current model parameters
 - 9: **Return** User and Item latent factor matrices U and V
-

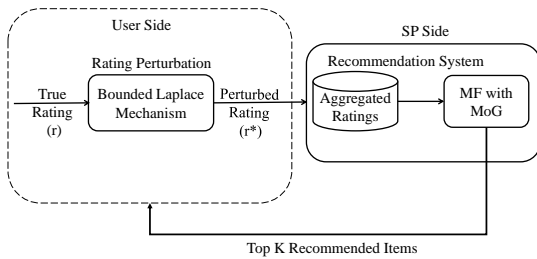


Figure 3: LDP based MF recommendation with MoG.

4.5 Predicted Ratings Combination

The predicted ratings combination module combines outputs from the CNN classification model and the MF-MoG model to produce the final predicted rating.

The final rating of a user an item is given as:

$$P_{final} = \beta * P_{MF} + (1 - \beta) * P_{Sent} \quad (1)$$

where P_{final} is the final predicted rating, P_{MF} is the rating predicted using the MF-MoG model, P_{Sent} is the rating predicted using sentiment analysis and β is the parameter that is used to adjust the importance of each component.

5 EVALUATION

In this section, we present the evaluation results of our proposed LDP based recommendation model. We use two datasets, Amazon Toys and Games and Amazon Instant Video, to validate the effectiveness of the proposed system. We use the Root Mean Squared Error (RMSE) to evaluate the prediction accuracy and F-score to evaluate the utility of the recommendation system.

5.1 Dataset

Table 1 provides a detailed view of the datasets we used in our evaluation.

Table 1: Datasets.

Dataset	Total Ratings	Total Reviews	Rating Scale
Amazon Toys and Games	2,252,771	167,597	1 to 5
Amazon Electronics	583,933	37,126	1 to 5

5.2 Metrics

The privacy budget ϵ acts as a metric of privacy loss at perturbed data. The lower the privacy budget ϵ is the higher the privacy that is guaranteed. We consider the value range from 0.1 to 3 for the privacy budget ϵ . The parameter β in Eq.(1) controls the roles MF-MoG and the sentiment analysis model play in determining the final predicted rating. The lower value of β indicates that the final predicted rating is more reliant on the sentiment analysis model, and the higher value means it relies more on the MF-MoG model. We evaluate the trade-off between β and utility by considering the value range from 0.1 to 1 for the parameter β . The utility of a recommendation system is evaluated by how well it predicts the relevance of an item for a user. We use F-score as the utility metric to measure

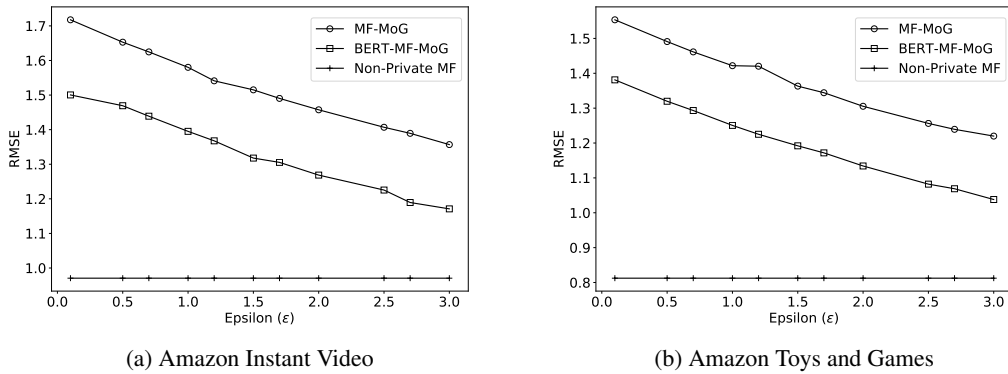


Figure 4: BERT-MF-MoG vs MF-MoG RMSE Comparison.

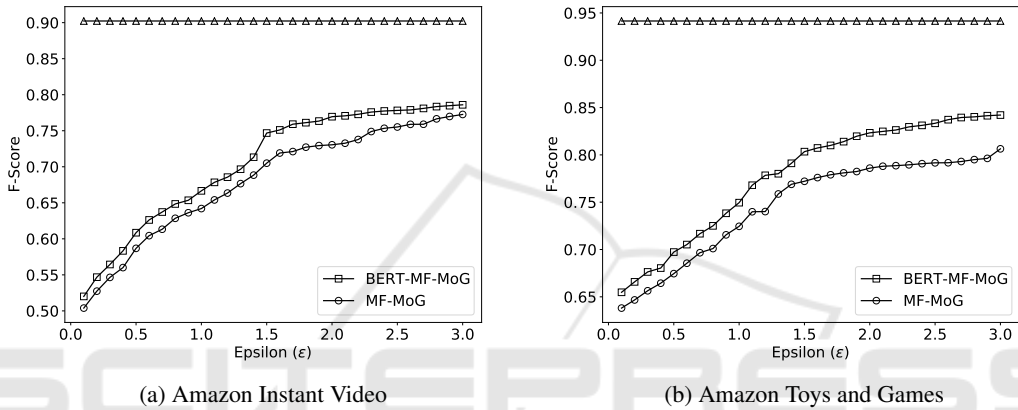


Figure 5: BERT-MF-MoG vs MF-MoG F-Score Comparison.

how well the recommendation systems make recommendations that adapt to a user’s choices. Then we use RMSE to measure the predictive accuracy of our system.

5.3 Results

In this experiment, we compare the prediction accuracy of our proposed recommendation model BERT-MF-MoG with MF-MoG (Neera et al., 2021) which, to the best of our knowledge, is the most comparable method as it also uses the exact input perturbation mechanism. We do not compare our method with other global or local differential privacy based recommendation systems due to the differences in perturbation mechanisms. The privacy budget ϵ varies from 0.1 to 3, and we use 0.7 as β values for both datasets. Fig. 4 and 4 shows the RMSE values for BERT-MF-MoG ($\beta = 0.7$), MF-MoG and the baseline method, Non-Private BERT-MF. The baseline method does not perturb the user’s original ratings, and the BERT review pre-processing takes place at the service provider’s side as no rating perturbation is performed, there is no need to combine MoG with MF. As ex-

pected, the prediction accuracy of the two privacy-preserving methods increases when the privacy budget ϵ increases. The results also show that BERT-MF-MoG produces a higher increase in recommendation accuracy for both datasets for all the values of ϵ than MF-MoG. This is because BERT-MF-MoG combines the sentiment classification with the MF-MoG model and predicted rating module takes into account both ratings and reviews. Fig.5 shows the F-score values for BERT-MF-MoG and MF-MoG models. The F-score value increases when the privacy budget ϵ increases for both methods. These F-score values demonstrate again that the BERT-MF-MoG model provides more accurate recommendations than MF-MoG for all privacy budget values ϵ .

Fig.6 and Fig.7 shows the RMSE and the F-score values respectively when varying β while keeping privacy budget at $\epsilon = 1.5$. These figures show that optimal RMSE and F-score values are obtained when $\beta = 0.7$. When $\beta = 0$, only the output of the sentiment analysis model, P_{Sent} , is used to determine the final predicted rating as indicated in Eq.(1), and the output of the MF-MoG model is not used. Only a few ratings are used as sentiment labels for the CNN




Figure 6: β -RMSE Comparison.




Figure 7: β -F-Score Comparison.

model, and they are perturbed at the user side. When β increases from 0, the contribution of the MF-MoG model, which uses a large proportion of ratings, starts to be included in the final predicted rating. Hence, a decrease in RMSE (increase in F-Score) can be observed as β increases in Fig.6 and Fig.7.

When β increases from 0.7 to 1, the RMSE value increases. This is because the prediction relies more heavily on the MF-MoG model than it should in this range. The significance of sentiment analysis is underestimated. When $\beta = 1$, Eq.(1) uses the P_{MF} as the final predicted rating. The output of the sentiment classification model P_{Sent} does not play any role in the prediction. The same trend can be observed in Fig.7 for the F-score value. The larger the β is, the more contribution the MF-MoG model makes, and the more accurate the prediction is until β reaches 0.7, where the highest F-score and lowest RMSE value are obtained. We use F-score as a utility metric and RMSE as the accuracy metric for our recommendation system. The F-score and RMSE results show that the BERT-MF-MoG model provides the most accurate recommendations when $\beta = 0.7$ for the two datasets. The optimal value of Beta should be obtained through numerical evaluation for different datasets.

6 CONCLUSION

In this paper, we proposed an LDP-based recommendation system that incorporates a deep-learning sentiment analysis model into a collaborative filtering algorithm. The system protects the privacy of users and at the same time offers substantial utility to the service provider. The recommendation accuracy of our proposed model is improved by taking advantage of sentiment analysis performed on user reviews. The experiments conducted with two amazon review and rating datasets demonstrated that the utility of our proposed recommendation system outperforms that of the recommendation system based just on ratings for all the values of privacy budget ϵ . In future work, we plan to explore using other techniques such as LSTM (Long Short Term Memory Networks) in combination with CNN for sentiment classification to further improve the recommendation accuracy.

REFERENCES

Cheng, Z., Ding, Y., Zhu, L., and Kankanhalli, M. (2018). Aspect-aware latent factor model: Rating prediction

- with ratings and reviews. In *Proceedings of the 2018 world wide web conference*, pages 639–648.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407.
- Goularas, D. and Kamis, S. (2019). Evaluation of deep learning techniques in sentiment analysis from twitter data. In *2019 International Conference on Deep Learning and Machine Learning in Emerging Applications (Deep-ML)*, pages 12–17. IEEE.
- Holohan, N., Antonatos, S., Braghin, S., and Mac Aonghusa, P. (2018). The bounded laplace mechanism in differential privacy. *arXiv preprint arXiv:1808.10410*.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. (2019). Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Li, X.-B. and Sarkar, S. (2011). Protecting privacy against record linkage disclosure: A bounded swapping approach for numeric data. *Information Systems Research*, 22(4):774–789.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pre-training approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Z., Wang, Y.-X., and Smola, A. (2015). Fast differentially private matrix factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 171–178.
- McSherry, F. and Mironov, I. (2009). 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*, pages 627–636.
- Meng, X., Wang, S., Shu, K., Li, J., Chen, B., Liu, H., and Zhang, Y. (2018). Personalized privacy-preserving social recommendation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Mobasher, B., Burke, R., Bhaumik, R., and Williams, C. (2007). Toward trustworthy recommender systems: An analysis of attack models and algorithm robustness. *ACM Transactions on Internet Technology (TOIT)*, 7(4):23–es.
- Neera, J., Chen, X., Aslam, N., Wang, K., and Shu, Z. (2021). Private and utility enhanced recommendations with local differential privacy and gaussian mixture model. *arXiv preprint arXiv:2102.13453*.
- Osman, N., Noah, S., and Darwich, M. (2019). Contextual sentiment based recommender system to provide recommendation in the electronic products domain. *International Journal of Machine Learning and Computing*, 9(4):425–431.
- Pal, A., Parhi, P., and Aggarwal, M. (2017). An improved content based collaborative filtering algorithm for movie recommendations. In *2017 tenth international conference on contemporary computing (IC3)*, pages 1–3. IEEE.
- Paterek, A. (2007). Improving regularized singular value decomposition for collaborative filtering. In *Proceedings of KDD cup and workshop*, volume 2007, pages 5–8.
- Raghavan, S., Gunasekar, S., and Ghosh, J. (2012). Review quality aware collaborative filtering. In *Proceedings of the sixth ACM conference on Recommender systems*, pages 123–130.
- Salinca, A. (2017). Convolutional neural networks for sentiment classification on business reviews. *arXiv preprint arXiv:1710.05978*.
- Shen, Y. and Jin, H. (2014). Privacy-preserving personalized recommendation: An instance-based approach via differential privacy. In *2014 IEEE International Conference on Data Mining*, pages 540–549. IEEE.
- Shin, H., Kim, S., Shin, J., and Xiao, X. (2018). Privacy enhanced matrix factorization for recommendation with local differential privacy. *IEEE Transactions on Knowledge and Data Engineering*, 30(9):1770–1782.
- Sutanto, J., Palme, E., Tan, C.-H., and Phang, C. W. (2013). Addressing the personalization-privacy paradox: An empirical assessment from a field experiment on smartphone users. *MIS quarterly*, pages 1141–1164.
- Wang, C., Duan, Q., Tong, C. H., Di, Z., and Gong, W. (2016). A gui platform for uncertainty quantification of complex dynamical models. *Environmental Modelling & Software*, 76:1–12.
- Wang, Y., Wang, M., and Xu, W. (2018). A sentiment-enhanced hybrid recommender system for movie recommendation: a big data analytics framework. *Wireless Communications and Mobile Computing*, 2018.
- Zarzour, H., Al-Ayyoub, M., Jararweh, Y., et al. (2021). Sentiment analysis based on deep learning methods for explainable recommendations with reviews. In *2021 12th International Conference on Information and Communication Systems (ICICS)*, pages 452–456. IEEE.