IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

"Enhanced Privacy-Preserving Recommendation Model Using Federated Learning"

Ranjeet Pratap Singh
M. Tech Scholar
CSE, Department of Computer Science Engineering,
NIIST Bhopal

Prof. Vaibhav Patel
Assistant Professor
CSE, Department of Computer Science Engineering,
NIIST Bhopal

Prof. Anurag Shrivastava
Assistant Professor
CSE, Department of Computer Science Engineering,
NIIST Bhopal

Abstract: In the era of digital personalization, recommendation systems play a crucial role in enhancing user experience by suggesting relevant content based on individual preferences. However, the centralized collection and processing of user data in traditional recommendation systems pose significant privacy risks. This research proposes a novel privacy-preserving recommendation model using Federated Learning (FL) to address these concerns. By decentralizing the training process, user data remains on local devices, and only model updates are shared with a central server, ensuring that sensitive information is never exposed. To further protect user privacy, we integrate Differential Privacy (DP) techniques at both the local training and model aggregation stages. An adaptive clipping strategy is also introduced to optimize the balance between privacy protection and recommendation accuracy. The proposed model is evaluated using real-world datasets, focusing on key metrics such as accuracy, privacy loss, and computational efficiency. Our results demonstrate that the model effectively safeguards user privacy while delivering high-quality recommendations. This approach offers a scalable and secure solution for privacy-preserving recommendation systems, paving the way for future advancements in privacy-centric applications across various industries.

Keywords— Federated Learning, Privacy Preservation, Differential Privacy, Recommender systems

I. INTRODUCTION

In today's digital age, recommendation systems have become integral to personalized user experiences, driving content discovery across various platforms such as e-commerce, streaming services, and social media. These systems analyze vast amounts of user data to provide tailored suggestions, enhancing user satisfaction and engagement. However, the effectiveness of traditional recommendation systems often comes at the cost of user privacy, as they rely on centralized data collection and processing. This centralization poses significant privacy risks, including unauthorized data access, data breaches, and misuse of personal information [5]. With growing concerns about data privacy and stringent regulations like GDPR (General Data Protection Regulation) and CCPA

(California Consumer Privacy Act), there is an urgent need for innovative approaches that can protect user privacy without compromising the quality of recommendations. Federated Learning (FL) has emerged as a promising solution to this challenge. FL enables the training of machine learning models across distributed devices, allowing data to remain local while only model updates are shared and aggregated at a central server. This decentralized approach significantly reduces the risk of data exposure and aligns with privacy-by-design principles [7]. While Federated Learning enhances privacy, it is not without its challenges. One key issue is the potential leakage of sensitive information through model updates, which can be exploited by malicious actors. To address this, Differential Privacy (DP) technique's can be integrated into the FL process, adding noise to the data or model updates to ensure that individual user information cannot be inferred from the shared models. This combination of FL and DP forms the foundation of our proposed enhanced privacy-preserving recommendation model.

This research aims to develop a robust recommendation system that leverages the strengths of Federated Learning and Differential Privacy to safeguard user data. By implementing adaptive privacy-preserving techniques, our model seeks to optimize the trade-off between privacy and recommendation accuracy. The proposed model is evaluated on real-world datasets to demonstrate its effectiveness in providing personalized recommendations while maintaining high standards of privacy protection, offering a scalable and secure solution for modern recommendation systems.

Federated Learning:

Federated Learning (FL) is an innovative approach to machine learning that enables the training of models across multiple decentralized devices or servers, such as smart phones, without requiring raw data to be shared with a central server. Instead of sending user data to a central location for processing, FL allows individual devices to train a local model using their own data. These local models are then sent to a central server, where they are aggregated to create a global model. This global model is subsequently shared back with the devices, which continue the iterative training process. The primary advantage of FL is its ability to maintain user privacy, as sensitive data never leaves the user's device. This decentralized approach reduces the risk of data breaches and aligns with privacy regulations like GDPR and CCPA. Additionally, FL can significantly reduce the need for data transfer, improving efficiency in environments with limited bandwidth.

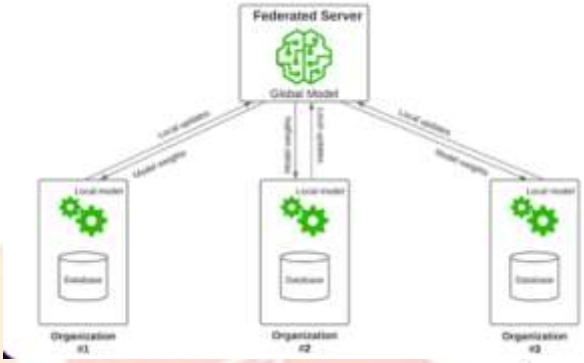

Figure 1.1 Federated Learning

Federated Learning is particularly valuable in scenarios where data is sensitive or distributed across many devices, such as in healthcare, finance, and personalized services. By keeping data local and only sharing model updates, FL offers a promising solution for building robust, privacy-preserving machine learning systems in the modern, data-driven world.

Privacy Preservation:

Privacy-preserving refers to the methods, techniques, and practices aimed at protecting the privacy of individuals or entities while processing, sharing, or storing their data. In the context of technology, particularly in machine learning and data analysis, privacy-preserving methods are designed to ensure that sensitive information is not exposed, leaked, or misused, even as data is utilized for beneficial purposes like model training or generating insights.

Federated Recommendation System:

A Federated Recommendation System is an advanced type of recommendation system that leverages Federated Learning (FL) to generate personalized recommendations while preserving user privacy. Traditional recommendation systems often require collecting large amounts of user data on a central server, raising privacy concerns. In contrast, a federated recommendation system allows user data to remain on local devices, with only model updates shared across the network, ensuring sensitive information is not exposed.

The paper also explores emerging trends such as hybrid approaches that combine multiple privacy-preserving techniques to leverage their respective strengths. Additionally, it discusses the application-specific challenges and solutions, highlighting case studies from various industries to illustrate practical implementations [5]. By providing a thorough examination of these strategies and their implications, this survey aims to bridge the gap between theoretical research and practical deployment, ensuring that federated learning can be safely and effectively used in privacy-sensitive environments. Ultimately, this survey serves as a valuable resource for advancing the field of privacy-preserving federated learning [6].

II. LITRETURE REVIEW

The literature survey provides an extensive review of existing research on privacy preservation in federated learning, examining a range of techniques such as differential privacy, secure multiparty computation, and homomorphic encryption. It highlights the strengths and limitations of these methods, offering insights into their practical applicability and performance. Additionally, the survey identifies emerging trends and key challenges, setting the stage for future advancements in the field.

Authors [1] proposed a Recommendation system (RecSys's) based on Deep Learning. This has resulted in the collection of substantial amounts of personal data on many platforms in recent years, leading to a data privacy problem. In this work FL has emerged as a technique that basically provides privacy and is therefore used in many scenarios where data privacy is of high priority. Consequently, we presented a movie RecSys, which is being trained end to- end using FL and scales well to exceptionally large numbers of users. This Proposed model helps to increase the accuracy.

Authors [2] dedicated to surveying of state-of-the-art privacy-preservation techniques in FL in relations with GDPR requirements. Furthermore, insights into the existing challenges are examined along with the prospective approaches following the GDPR regulatory guidelines that FL-based systems shall implement to fully comply with the GDPR.

Authors [3] proposed a efficient privacy-preserving recommender system based on user classification, RSUC uses Paillier encryption to enable secure computation for protecting user privacy. RSUC takes advantage of extra attributes collected by service providers such as user attributes to classify users before computing user similarity and recommendation, reducing the amount of data needed for computation.

Author [4] explored how federated recommendation systems allow for secure data sharing and collaboration in the context of big data. In addition, we identified the challenges associated with developing and deploying these systems in the real world. With the increasing demand for privacy, federated recommendation systems provide a powerful tool for protecting user data while allowing organizations to benefit from the data they share.

Author [5] proposed a framework functions with arbitrary types of input features that emphasize its usability with natural language data. The text input on the client-side is encoded using a rolling hash-based representation, which provides a combined

solution for the high resource demands of embedding algorithms and the privacy concerns of sharing sensitive data. Authors evaluate method in a sentiment analysis task using the IMDB Movie Reviews dataset as well as a rating prediction task with the Movie Lens dataset augmented with additional movie keywords.

In this work [6] is dedicated to surveying of state-of-the-art privacy-preservation techniques in FL in relations with GDPR requirements. Furthermore, insights into the existing challenges are examined along with the prospective approaches following the GDPR regulatory guidelines that FL-based systems shall implement to fully comply with the GDPR.

In [7] result suggests that proposed algorithm is an effective method of implementing differential privacy with federated learning, and clinical data scientists can use our general framework to produce differentially private models on federated datasets.

In [8] conduct a detailed study on FL, the categorization of FL, the challenges of FL, and various attacks that can be executed to disclose the users' sensitive data used during learning. In this survey, authors review and compare different privacy solutions for FL to prevent data leakage and discuss secret sharing (SS)-based security solutions for FL proposed by various researchers in concise form. Authors also briefly discuss quantum federated learning (QFL) and privacy-preservation techniques in QFL.

In this [9] paper, we reiterate the concept of federated learning and propose secure federated learning (SFL), where the ultimate goal is to build trustworthy and safe AI with strong privacy-preserving and IP-right-preserving. We provide a com- prehensive overview of existing works, including threats, attacks, and defenses in each phase of SFL from the lifecycle perspective.

III. PROPOSED RECOMMONDATION SYSTEM

The proposed methodology for the Enhanced Privacy-Preserving Recommendation Model using Federated Learning focuses on developing a secure and efficient recommendation system that protects user privacy. The key steps are:

Federated Learning Setup:

Data Distribution: The model will be trained on user interaction data (e.g., movie ratings, product clicks) distributed across multiple clients (e.g., mobile devices).

Local Model Training: Each client will train a recommendation model (e.g., Collaborative Filtering) on its local data. The training will include user-item interaction matrices, which are updated locally.

Model Aggregation: The central server will aggregate the locally trained models using Federated Averaging (FedAvg), combining model updates without accessing raw user data.

Differential Privacy Integration

DP-SGD (Differentially Private Stochastic Gradient Descent): Implement DP-SGD at the client level to add noise to the gradient updates, ensuring individual user data cannot be reverse-engineered from the updates.

Privacy-Preserving Aggregation: Introduce noise during the aggregation process on the server side to further protect the combined model from leaking sensitive information.

Adaptive Clipping: Apply adaptive clipping strategies to limit the influence of any single data point before adding noise, optimizing the trade-off between privacy protection and model accuracy.

Model Architecture

Deep Learning-Based Recommendation Model: Each user's device hosts a local deep learning model designed for generating recommendations. Common architectures used in this context include:

Neural Collaborative Filtering (NCF): This model combines the strengths of matrix factorization and deep neural networks. It uses embeddings for users and items, which are fed into a multi-layer perceptron (MLP) to learn complex interaction patterns between them.

Autoencoders: For content-based filtering, autoencoders can be used to learn latent representations of items based on their features and user interactions, helping to predict user preferences.

Local Training Process: Implement the model architecture on each client, using local user data for training, and periodically send the model updates to the central server.

Evaluation Metrics

Accuracy Metrics: Measure the performance of the recommendation model using Precision, Recall, F1-score, and Mean Squared Error (MSE).

Privacy Metrics: Evaluate the privacy guarantees using the privacy loss parameter (ε) and the probability of privacy failure (δ).

Efficiency Metrics: Analyze training time, communication overhead, and scalability to assess the efficiency of the federated learning approach

IV. EXPERIMENTAL SETUP & RESULT ANALYSIS

Federated Recommendation System, choosing the right dataset is crucial for evaluating the system's performance in terms of both recommendation accuracy and privacy preservation. In this work MovieLens and Amazon product review dataset are used for the experimental purpose.

MovieLens : MovieLens datasets are widely used for benchmarking recommendation systems. They consist of movie ratings provided by users, along with metadata like movie titles, genres, and user demographics. This dataset is ideal for testing collaborative filtering algorithms and deep learning models in a federated setting. The decentralized nature of user ratings in federated learning aligns well with the dataset's structure. This is available in various sizes, from 100,000 ratings (MovieLens 100K) to 26 million ratings (MovieLens 25M). By keeping the user ratings on local devices, federated learning can be tested for privacy preservation while still offering personalized movie recommendations.

Amazon Product Reviews: This dataset contains product reviews and ratings from Amazon users across various categories like electronics, books, and clothing. It includes metadata such as product descriptions, prices, and user profiles. Useful for developing a federated recommendation system for e-commerce, this dataset can help in building models that recommend products based on users' past purchasing behavior. The dataset contains millions of reviews and ratings, making it suitable for large-scale federated learning experiments. The dataset can be used to simulate a federated environment where users' purchase history is kept private on their devices, with only model updates being shared.

Python tool and libraries like TensorFlow Federated, PySyft are used for the experimental purpose. Performance of proposed model on different dataset shown on the table 4.1

Table 4.1 Performance of Proposed Model on different dataset

Dataset	Model	RMSE	MAE	Accuracy
MovieLens	Autoencoder + DP-SGD + FedAvg	0.86	0.68	88%
Amazon Review	NCF + DP-SGD + FedAvg	0.93	0.72	71%

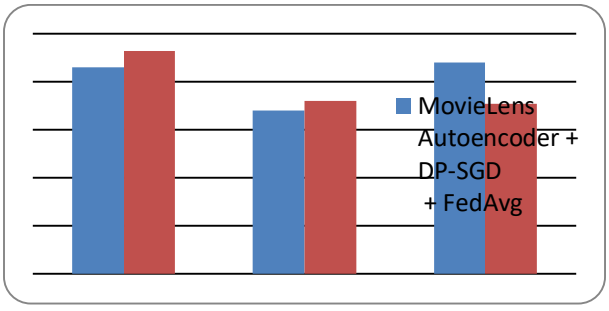

Figure 4.1 Performance Graph of Proposed Model

CONCLUSION

This research work represents a significant advancement in safeguarding user privacy while delivering personalized recommendations. By leveraging Federated Learning, the model ensures that user data remains decentralized, drastically reducing the risk of data breaches and unauthorized access. The integration of Differential Privacy techniques further strengthens privacy protections by preventing the leakage of sensitive information through model updates. Our evaluation of the model on real-world datasets demonstrates its ability to maintain high levels of recommendation accuracy while upholding stringent privacy standards. The adaptive privacy-preserving strategies employed in this model effectively balance the trade-off between privacy and performance, ensuring that users can enjoy tailored content recommendations without compromising their personal data. This model offers a scalable, secure solution for modern recommendation systems, making it highly relevant for applications in industries where privacy is paramount. It sets a foundation for future work in developing even more robust privacy-preserving machine learning technologies.

REFERENCES

[1] David Neumann, et al "A Privacy Preserving System for Movie Recommendations Using Federated Learning" ACM Trans. Recomm. Syst., Vol. 1, No. 1, Article 1. Publication date: January 2023

[2] Nguyen Truong et al "Privacy preservation in federated learning: An insightful survey from the GDPR perspective" Published by Elsevier Ltd, 2021

[3] Junwei Luo et al "Privacy-preserving recommendation system based on user classification" https://doi.org/10.1016/j.jisa.2023.103630

[4] Muhammad Asad et al "A Comprehensive Survey on Privacy-Preserving Techniques in Federated Recommendation Systems" Appl. Sci. 2023, 13, 6201. https://doi.org/10.3390/app13106201

[5] Balázs Nagy et al." Privacy-preserving Federated Learning and its application to natural language processing" https://doi.org/10.1016/j.knosys.2023.110475 Published by Elsevier 2023

[6] Nguyen Truong et. al. "Privacy preservation in federated learning: An insightful survey from the GDPR perspective" https:// doi.org/10.1016/j.cose.2021.102402 Elsevier Ltd 2021

[7] Amol Khanna et al "Privacy-preserving Model Training for Disease Prediction Using Federated Learning with Differential Privacy" 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) Scottish Event Campus, Glasgow, UK, July 11-15, 2022

[8] Sanchita Saha1, "A multifaceted survey on privacy preservation of federated learning: progress, challenges, and opportunities" Artificial Intelligence Review (2024) 57:184 https://doi.org/10.1007/s10462-024-10766-7, 2024

[9] Qiang Yang et al "Federated Learning with Privacy-preserving and Model IP-right-protection" 20(1), February 2023, 19-37 DOI: 10.1007/s11633-022-1343-2 www.mi-research.net

[10] Huiyong Wang1, et. al. "Privacy-preserving federated learning based on partial low-quality data" Wang et al. Journal of Cloud Computing (2024) 13:62 https://doi.org/10.1186/s13677-024-00618-8

[11] Georgios A. Kaissis et. al. "Secure, privacy-preserving and federated machine learning in medical imaging" Nature Machine Intelligence | VOL 2 | June 2020 | 305–311 | www.nature.com/natmachintell

[12] Q. Yang, Y. Liu, T. Chen, Y. Tong, Federated machine learning: Concept and applications, ACM Trans. Intell. Syst. Technol. 10 (2) (2019) 1–19, http://dx.doi.org/10.1145/3298981.

[13] Q. Xia, W. Ye, Z. Tao, J. Wu, Q. Li, A survey of federated learning for edge computing: Research problems and solutions,

High-Confi. Comput. (2021) 100008, http://dx.doi.org/10.1016/j.hcc.2021.100008.

[14] M. Aledhari, R. Razzak, R.M. Parizi, F. Saeed, Federated learning: A survey on enabling technologies, protocols, and applications, IEEE Access 8 (2020) 140699–140725, http://dx.doi.org/10.1109/ACCESS.2020.3013541.

[15] Z. Li, Z. Huang, C. Chen, C. Hong, Quantification of the leakage in federated learning, 2020, arXiv:1910.05467.

[16] A. Wainakh, F. Ventola, T.M. ig, J. Keim, C.G. Cordero, E. Zimmer, T. Grube, K. Kersting, M. Mühlhäuser, User label leakage from gradients in federated learning, 2021, arXiv:2105.09369.

[17] S.P. Kasiviswanathan, H.K. Lee, K. Nissim, S. Raskhodnikova, A. Smith, What can we learn privately? SIAM J. Comput. 40 (3) (2011) 793–826, http://dx.doi.org/10.1137/090756090.

[18] X. Xiong, S. Liu, D. Li, Z. Cai, X. Niu, A comprehensive survey on local differential privacy, in: A.M. Del Rey (Ed.), Secur. Commun. Netw. 2020 1–29, http://dx.doi.org/10.1155/2020/8829523.

[19] L. Sun, J. Qian, X. Chen, P.S. Yu, LDP-FL: Practical private aggregation in federated learning with local differential privacy, 2020, arXiv: 2007 15789.