Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 3, No. 8, 2024

Collaborative Optimization in Federated Recommendation: Integrating User Interests and Differential Privacy

Lipeng Zhu[1], Wanyu Cui[2], Yue Xing[3], Yang Wang[4]

[1]Johns Hopkins University, Baltimore, USA

[2]University of Southern California, Los Angeles, USA

[3]University of Pennsylvania, Philadelphia, USA

[4]University of Michigan, Ann Arbor, USA

*Corresponding Author: Yang Wang; oceanwy@umich.edu

Abstract: This paper addresses the conflict between personalization effectiveness and privacy protection in federated recommendation systems. It proposes a collaborative optimization method that integrates local interest guidance with a differential privacy mechanism. The goal is to enhance both recommendation performance and security under the condition of data isolation across multiple parties. Specifically, a personalized model aggregation strategy based on local interest embeddings is designed. By incorporating user preference features into the global model update process, the model can adaptively capture individual differences among clients during aggregation. At the same time, to reduce the negative impact of privacy protection on model performance, a differential privacy-driven personalized update mechanism is introduced. This mechanism ensures the non-inferability of user data while applying a gradient-guided noise regulation strategy. It helps preserve the local model's ability to represent individual interests. Multiple comparative experiments conducted on standard recommendation datasets show that the proposed method outperforms representative federated recommendation models across various metrics. It also demonstrates strong robustness and stability under highly non-independent data distributions and high noise settings. Further ablation studies confirm the independent and joint contributions of the two key modules in enhancing the model's personalization capacity and resistance to interference. These results highlight the method's ability to achieve an effective balance between privacy protection and recommendation quality.

Keywords: Federated recommendation, personalized modeling, differential privacy, interest-guided aggregation

1. Introduction

In the current digital age, the explosion of information and the flood of data have made users increasingly reliant on recommendation systems to access personalized content. These systems analyze user behavior, preference patterns, and content features to achieve accurate individual-level information matching. However, traditional recommendation systems are often built upon the centralized collection and processing of large-scale data [1,2]. While this architecture improves recommendation accuracy to some extent, it also leads to serious privacy concerns. As users become more aware of data security, protecting personal privacy has become a fundamental requirement for building trustworthy recommendation systems. Therefore, achieving effective personalized recommendations while preserving user privacy has become one of the key challenges in recommendation system research [3].

Federated learning, as an emerging paradigm of distributed machine learning, enables collaborative model training without sharing raw data. Its core idea is to keep data local while multiple participants jointly train a global model. This approach helps mitigate privacy leakage risks. Integrating federated learning into recommendation systems offers a novel solution to the trade-off between privacy protection and personalized modeling. With this mechanism, recommendation systems can support cross-device and cross-platform information sharing without centralizing users' sensitive data. This not only enhances the flexibility of data usage but also builds user trust in the platform, laying a technical foundation for the continued development of privacy-aware recommendation systems [4].

However, federated recommendation systems also face significant challenges. Data heterogeneity across user devices, imbalanced computational capabilities, and limited communication bandwidth all place greater demands on model training stability and efficiency. More critically, privacy-preserving mechanisms such as differential privacy or secure multi-party computation often reduce model performance while protecting data. The core value of recommendation systems lies in accurately predicting user interests [5,6]. Excessive privacy constraints may weaken personalized modeling. Thus, achieving coordinated optimization between privacy protection and model accuracy is crucial for making federated recommendation systems practical [7].

From a modeling perspective, personalized recommendation is essentially a dynamic learning process. It requires continuously capturing changes in user interests and

contextual features. In a federated environment, personalized modeling demands models with strong generalization ability and fast adaptation to individual preferences based on local data. This process is influenced by various factors, such as data sparsity, limited feedback signals, and the non-stationary nature of user behavior [8]. To enhance personalization, it is necessary to design more efficient user modeling strategies. These strategies must be deeply integrated with the federated learning mechanism to ensure that recommendation models can respond to individual differences while maintaining overall learning stability and efficiency.

Against this backdrop, studying the coordinated optimization of privacy protection and personalized modeling in federated recommendation systems has both theoretical and practical significance. On one hand, it promotes the application of privacy-preserving technologies in recommendation scenarios, offering a more balanced technical solution between data security and intelligent services. On the other hand, such research supports the evolution of recommendation system architectures toward more distributed and transparent frameworks, laying the groundwork for multi-party collaborative intelligence [10]. Furthermore, these findings can provide general-purpose paradigms and solutions for personalized services in other sensitive domains, such as smart healthcare, financial risk control, and social media. In summary, the coordinated optimization of privacy protection and personalization in federated recommendation systems not only responds to both technological progress and societal needs but also offers a forward-looking direction for the development of trustworthy artificial intelligence.

2. Related work

2.1 Federated Learning

Federated learning, as a rapidly evolving paradigm of distributed machine learning, aims to enable collaborative modeling across multiple parties while preserving data privacy. Its core idea is to train models locally on multiple clients without sharing raw data [11]. These clients upload local model parameters or gradients to a central server for aggregation and global model updating [12]. This process is typically carried out in multiple rounds. It not only reduces the risk of data leakage but also avoids single points of failure and storage bottlenecks that may occur in traditional centralized learning. Federated learning was originally designed to address the problem of data silos in sensitive domains such as finance and healthcare. In these sectors, data are difficult to share but high-quality joint modeling is still needed. Therefore, federated learning safeguards data sovereignty while providing a new path for large-scale machine learning under privacy constraints [13].

Key topics in federated learning research include communication efficiency, model aggregation strategies, and handling system heterogeneity. Since participating devices often vary in computing power, network conditions, and data distributions, training can suffer from inefficiencies or degraded model performance. To address this, researchers have proposed various communication optimization techniques. These include model compression, update frequency control, and asynchronous training mechanisms to reduce communication overhead[14]. Meanwhile, to improve performance under non-independent and identically distributed (Non-IID) data, many aggregation strategies have been enhanced. Examples include personalized modeling, weighted aggregation, and adversarial training. These methods improve the global model's ability to adapt to local differences. Such developments significantly increase the practicality and flexibility of federated learning in real-world applications [15].

Beyond the general federated learning framework, several task-specific variants have emerged, such as federated transfer learning, federated multi-task learning, and federated personalized learning. These variants demonstrate the adaptability of federated learning to different data properties and task requirements. For instance, in recommendation systems where user preferences are highly diverse, federated personalized learning focuses on improving individual model performance without compromising user privacy. This is often achieved through meta-learning, model hierarchies, or local fine-tuning. These techniques help balance local adaptation and global collaboration. They also enhance federated learning's ability to capture individual differences and support the shift from unified models to customized models. This better meets the needs of personalized services.

Overall, federated learning serves as a critical link between privacy protection and distributed collaboration. It shows great promise in both theoretical exploration and engineering practice[16]. Its initial applications in recommendation systems, speech recognition, and medical diagnosis have validated its potential as a new learning paradigm. However, federated learning still faces challenges related to stability, scalability, and security. These include preventing inference attacks during model synchronization and managing data updates from long-term offline devices. Therefore, further research on federated learning is not only important for advancing privacy-preserving machine learning but also essential for building trustworthy and reliable intelligent systems.

2.2 Personalized modeling recommendation algorithm

Personalized modeling is a central research topic in recommendation systems. Its goal is to generate content recommendations that match individual needs based on user behavior, interests, and contextual information [17]. As the number of users and the scale of information grow rapidly, recommendation systems are shifting from unified approaches to more fine-grained and adaptive personalization. The key to personalized modeling lies in accurately capturing user preferences and dynamically adjusting recommendation strategies to reflect changes in user interests over time. This process involves extracting effective features from user interactions, building robust user representations, and designing flexible model architectures. In recent years, advances in deep learning have supported the integration of multi-dimensional features such as historical behavior, content semantics, and

social relations. These developments have provided powerful tools for improving recommendation performance [18].

Currently, mainstream approaches to personalized recommendation can be categorized into collaborative filtering models, content-based models, and hybrid modeling strategies. Collaborative filtering identifies similarities between users or between users and items to make recommendations. It offers scalability and flexibility, and performs well in the absence of explicit feedback. However, it is sensitive to data sparsity and cold-start problems [19]. Content-based models rely on item attributes or user profiles to make predictions. They can partially address cold-start issues but often lead to information isolation and fail to capture user-to-user relationships. Hybrid modeling strategies have emerged as a promising direction. By combining collaborative filtering with content features and incorporating context-aware mechanisms, these methods enhance the generalization and diversity of recommendation systems [20].

A major challenge in personalized modeling lies in effectively managing the dynamic and nonlinear nature of user behavior. User preferences are rarely static; they change over time due to shifts in interests, contexts, or external influences. Traditional recommendation models often rely on static representations of long-term user preferences, which fail to capture these evolving patterns. This limitation reduces the ability of such models to provide timely and contextually relevant recommendations. As user engagement with content becomes increasingly fluid, the need for models that can dynamically adapt to behavioral changes becomes critical for delivering effective personalization [21].

To address this challenge, researchers have explored a range of advanced techniques, including sequential modeling, graph neural networks (GNNs), and attention mechanisms. Sequential modeling captures the temporal dependencies in user interactions, enabling the system to understand how recent behaviors influence current interests. Graph neural networks allow for the modeling of intricate relationships between users and items while maintaining structured representations of social or semantic information. Attention mechanisms further enhance modeling capacity by dynamically weighting behavioral features based on their relevance, enabling the model to focus on the most informative aspects of user activity. Together, these techniques expand the expressive power and adaptability of recommendation models, playing a key role in the ongoing advancement of personalized modeling approaches.

Despite the progress, personalized modeling still faces several practical challenges. First, it is essential to balance recommendation accuracy with model complexity and computational cost, especially in real-world deployments. Second, personalized modeling relies on large volumes of user data. This creates difficulties under increasingly strict data privacy regulations. Enabling effective personalization while preserving privacy has become a critical research problem. In addition, model interpretability is gaining attention, particularly in domains like finance and healthcare where transparency is crucial. The explainability of recommendations affects user trust and acceptance. Therefore, the development of personalized modeling must not only push technical boundaries but also achieve deeper integration and optimization in terms of privacy, security, and fairness.

3. Method

This study proposes a collaborative optimization approach for privacy protection and personalized modeling in federated recommendation systems. The goal is to improve recommendation performance and enhance individual adaptability without sharing data across parties. The method integrates federated learning architecture with personalized modeling strategies to deliver differentiated recommendations while preserving user privacy. First, a Local Interest-guided Personalized Model aggregation mechanism (LIPM) is introduced. It guides the federated aggregation process using user preferences learned locally on each client. This allows the global model to balance global optimization with local adaptation. As a result, the model becomes more robust to heterogeneous user data. Second, a Differential Privacy-driven Personalized Strategy (DPPS) is designed. It incorporates a dynamic perturbation mechanism during model updates. This strategy controls the risk of privacy leakage while maintaining high-fidelity modeling of individual interest patterns. It improves the trade-off between recommendation accuracy and privacy protection. These two innovations are complementary in both structural design and optimization objectives. Together, they promote the joint improvement of personalization and privacy in federated recommendation systems. The model architecture is shown in Figure 1.

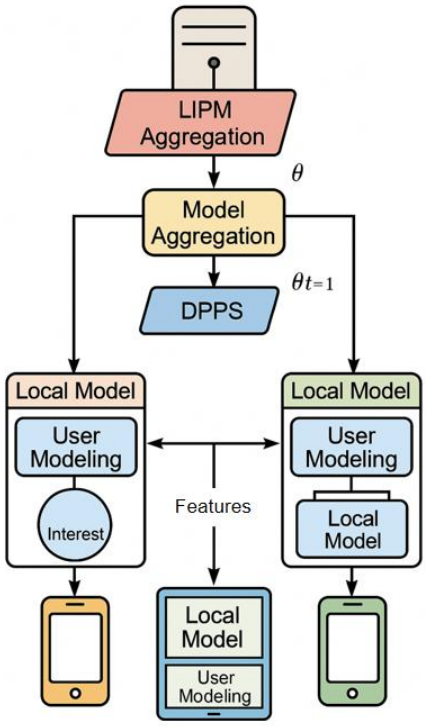

Figure 1. Overall model architecture diagram

3.1 Local Interest-guided Personalized Model

In federated recommendation systems, user data typically exhibit non-independent and identically distributed (Non-IID) properties due to varying user behaviors, preferences, and contexts. These differences lead to strong personalized needs across clients. Under such conditions, traditional model aggregation strategies that rely on uniform parameter averaging often fail to account for individual user characteristics. As a result, the global model struggles to generalize effectively across heterogeneous clients. The lack of personalization in the aggregation process limits the model's ability to deliver accurate and relevant recommendations, particularly in environments where user interests are diverse and dynamic.

To address this limitation, this paper introduces a personalized model aggregation mechanism guided by local interest representations. The core idea is to incorporate user-specific preference signals into the aggregation process, allowing the global model to better align with local user characteristics. By leveraging locally learned interest patterns, the proposed method enhances the adaptability of the federated model to personalized data distributions. This not only improves the model's performance on each client but also preserves the benefits of global collaboration. The module architecture of the proposed mechanism is illustrated in Figure 2.

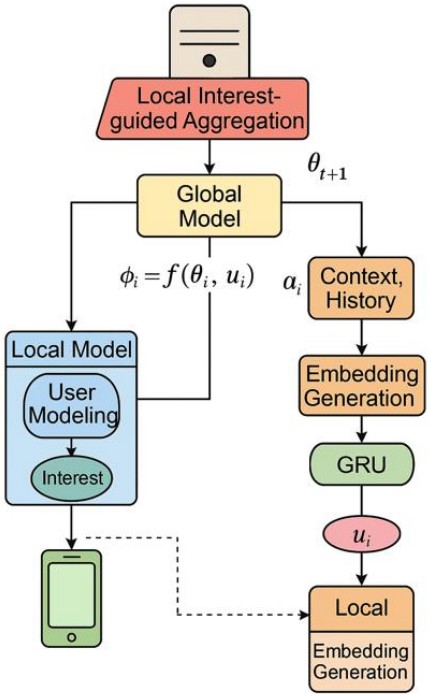

Figure 2. LIPM module architecture

In each round of communication, this method not only uploads local model parameters, but also introduces user interest representation as an auxiliary modeling signal to improve the personalized guidance of the aggregation strategy. Assuming that the local model of the i-th client is θ_i and its interest embedding is u_i, the comprehensive representation uploaded by the client can be defined as:

$$\phi_i = f(\theta_i, u_i)$$

Where $f(\cdot)$ represents the fusion function, which is used to capture the interaction information between model parameters and features of interest.

On the server side, in order to balance individual differences and overall performance in the global model aggregation process, a weighted aggregation strategy is adopted, and user interest guidance is embedded in the construction of aggregation weights. Define the global model as θ, then the update formula is:

$$\theta^{(t+1)} = \sum_{i=1}^{N} \alpha_i \cdot \phi_i$$

$\alpha_i = \dfrac{\exp(g(u_i))}{\sum_j \exp(g(u_j))}$ represents the attention weight calculated based on interest embedding, and $g(\cdot)$ is the weight generation function, which is used to measure the influence of each client interest on the global model. This design enables the aggregation mechanism to adaptively focus on user preferences that are more valuable to global performance improvement when facing heterogeneous data.

In addition, in order to further improve the expressiveness of local interest embedding, an embedding generation network combining contextual behavior and short-term history is adopted. Specifically, for user i's interest state at time t, represented as $u_i^{(t)}$, its update process can be modeled as:

$$u_i^{(t)} = GRU(x_i^{(t)}, u_i^{(t-1)})$$

Where $x_i^{(t)}$ is the feature representation of the current behavior, and GRU represents the gated recurrent unit structure, which is used to capture dependencies in time series. This structure can dynamically model the evolution of interests, thereby providing more expressive preference guidance signals for personalized aggregation.

Finally, in order to achieve synchronous update and differentiated adaptation of personalized models, a local fine-tuning strategy is designed after each round of communication. After receiving the aggregated model $\theta^{(t+1)}$, the client will make local adjustments to the model based on its interest representation u_i to make it more suitable for local preferences. The update strategy is as follows:

$$\theta_i^{(t+1)} = \theta^{(t+1)} + \lambda \nabla_\theta L_i(\theta^{(t+1)}, u_i)$$

Among them, λ is the adjustment coefficient, and L_i is the local personalized loss function, which comprehensively considers the matching degree between recommendation accuracy and interest expression. Through the above design, the local interest guidance mechanism not only improves the

expression of personalized information in the aggregation stage, but also achieves further alignment between the model and individual preferences through fine-tuning, thereby significantly enhancing the personalized ability of the recommendation system while protecting privacy.

3.2 Differential Privacy-driven Personalized Strategy

In order to further enhance the privacy protection capability in the federated recommendation system while maintaining the personalized expression capability of the model, this paper proposes a differential privacy-driven personalized strategy (DPPS). This method introduces a dynamic perturbation mechanism in the process of local model update and upload to control the risk of information leakage while maintaining the modeling accuracy of user preference features. Its module architecture is shown in Figure 3.

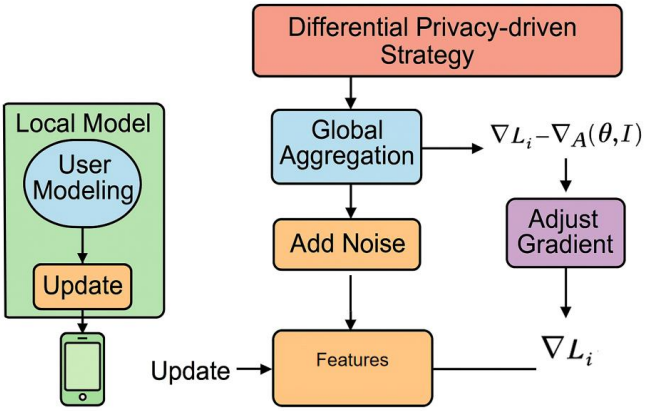

Figure 3. DPPS module architecture

Specifically, before uploading the model parameters, the client protects sensitive information by injecting differential privacy noise. For the i-th client, its local model parameter update is represented by $\theta_i^{(t)}$, and the perturbation version it uploads is:

$$\widetilde{\theta}_i^{(t)} = \theta_i^{(t)} + N(0, \sigma^2 I)$$

Where $N(0, \sigma^2 I)$ represents Gaussian noise with mean 0 and covariance $\sigma^2 I$, and the noise intensity σ controls the degree of privacy protection.

In order to maintain the accuracy of recommendations while protecting privacy, DPPS introduces a privacy budget control mechanism to dynamically adjust the noise size according to the user's behavior activity and historical disturbance effects. Let the privacy budget be ε_i, then the standard deviation of the Gaussian mechanism satisfies the relationship with it:

$$\sigma_i = \frac{\Delta}{\varepsilon_i}$$

Where Δ is the upper bound of sensitivity, which indicates the maximum change that model parameters may have when a single user's data changes. DPPS allocates personalized privacy budgets based on the activity of each client, thereby achieving differentiated protection of sensitive data.

In addition, considering that the noise introduced by differential privacy will interfere with the model training process, DPPS designs a personalized adjustment function to reversely guide the model gradient after adding noise, thereby enhancing the model's ability to respond to valid signals. Let the local loss function be $L_i(\theta)$ and its gradient be ∇L_i. The guided gradient adjustment is expressed as:

$$\widetilde{\nabla} L_i = \nabla L_i - \gamma \cdot \nabla N(0, \sigma^2 I)$$

γ is a regulatory factor, which is used to balance the learning signal and noise deviation and improve the model's ability to capture the preferred direction.

Finally, when performing global aggregation on the server side, DPPS maintains the standard weighted aggregation framework, but the aggregation weights will comprehensively consider the impact of noise introduction on model performance. The global model update is expressed as:

$$\theta^{(t+1)} = \sum_{i=1}^{N} w_i \cdot \widetilde{\theta}_i^{(t)}$$

Where $w_i = \frac{1}{\sigma_i^2} / \sum_{j=1}^{N} \frac{1}{\sigma_j^2}$ represents a weighting strategy based on the inverse of the privacy noise variance, which makes clients with less noise and higher stability have a greater impact on the global model, thereby effectively alleviating performance fluctuations caused by privacy disturbances. Through the above mechanism, DPPS retains the modeling ability of local preferences while achieving the synergistic goal of personalized modeling and privacy control while realizing the privacy protection of federated learning.

4. Experimental Results

4.1 Dataset

This study uses the real-world dataset MovieLens-1M, which is widely adopted in the recommendation system domain for evaluating models and validating algorithms. MovieLens-1M contains approximately one million user rating records. It covers around 6,000 users and 4,000 movies. Each record includes a user ID, movie ID, rating value, and timestamp. This information reflects long-term user preferences and behavioral patterns.

The dataset features a wide distribution of user activity and diverse rating behaviors. It is suitable for studying recommendation modeling under non-independent and identically distributed (Non-IID) conditions. It also supports scenarios involving federated learning and privacy protection. The historical rating sequences allow time-series modeling, which helps capture the evolution of user interests.

During data preprocessing, ratings are usually binarized. For example, ratings of 4 or higher may be treated as positive

feedback. The data are split into training and test sets based on users, simulating personalized inference tasks in real recommendation scenarios. With the characteristics of MovieLens-1M, this study can systematically evaluate the model's personalized recommendation performance under privacy-preserving conditions while ensuring data realism.

4.2 Experimental setup

In the experimental setup, we use a federated learning framework to simulate a scenario where multiple clients participate in the recommendation task. Each client holds a distinct subset of users, forming a typical non-independent and identically distributed data partition. This design reflects a realistic data silo environment under privacy constraints. The model is trained locally on each client and aggregated globally on the server. A differential privacy mechanism is introduced during training to control the sensitivity of uploaded information. A personalization strategy based on interest modeling enhances local adaptability.

We evaluate the approach using the MovieLens-1M dataset. The model architecture is a multi-layer perceptron (MLP). Key parameters, including embedding dimension, learning rate, and privacy budget, are optimized using grid search. This ensures that each method is compared under its best configuration. Evaluation metrics include AUC and NDCG, which measure the model's ranking performance and the effectiveness of personalized recommendation, respectively. The main experimental hyperparameter settings are shown in Table 1.

Table 1: Hyperparameter setting

Parameter	Value
Dataset	MovieLens-1M
Embedding Dimension	64
Learning Rate	0.001
Batch Size	128
Local Epochs	5
Communication Rounds	50
Privacy Budget	{0.5, 1.0, 5.0}
Optimizer	Adam

4.3 Experimental Results

1) Comparative experimental results

First, this paper gives the comparative experimental results with other models. The experimental results are shown in Table 2.

Table 2: Comparative experimental results

Method	NDCG@10	Recall@10	Precision@10
FedRec[22]	0.410	0.532	0.326
FedMF[23]	0.428	0.548	0.339
FedNCF[24]	0.441	0.562	0.351
DP-FedAvg[25]	0.417	0.529	0.323
PFedMe[26]	0.453	0.571	0.357
Ours	0.481	0.595	0.374

As shown in the table, the proposed method outperforms existing mainstream federated recommendation approaches across all three evaluation metrics: NDCG@10, Recall@10, and Precision@10. Specifically, NDCG@10 improves from 0.410 in FedRec to 0.481. This indicates that the model has a stronger ability to rank items that users are interested in. The improvement in this metric is especially important, as ranking relevant items higher directly affects user experience and satisfaction in real-world recommendation systems.

From the perspective of Recall@10, the proposed method achieves a recall of 0.595. This is significantly higher than that of advanced methods such as FedNCF (0.562) and pFedMe (0.571). This result suggests that the method can capture a larger portion of truly relevant items within a limited recommendation list. It shows stronger adaptability to user preference diversity and improves the comprehensiveness of recommendations. This is particularly valuable in federated learning settings, where data heterogeneity is common. A model's ability to provide personalized recommendations is crucial for its effectiveness in such environments.

For the Precision@10 metric, the proposed method also demonstrates higher accuracy. It increases from 0.323 in DP-FedAvg to 0.374. This means that the top 10 items in the recommendation list include more truly relevant content. It reduces the presence of irrelevant items and enhances user trust and efficiency in interacting with the recommendations. When combined with the previous two metrics, it is clear that the method achieves a balance between recommendation breadth and accuracy. This is a strength that many existing approaches struggle to maintain.

Taken together, the performance on all three metrics shows that the proposed method effectively addresses the performance degradation commonly seen in federated recommendation systems. It achieves this under privacy-preserving conditions by combining local interest guidance with a differential privacy strategy. The method improves the model's ability to capture personalized preferences while ensuring data security. This confirms its potential value and applicability in practical scenarios.

2) Ablation Experiment Results

Secondly, this paper gives the ablation experiment results, as shown in Table 3.

Table 3: Ablation Experiment Results

Method	NDCG@10	Recall@10	Precision@10
BaseLine	0.436	0.556	0.344
+LIPM	0.462	0.578	0.361
+DPPS	0.451	0.567	0.353
Ours	0.481	0.595	0.374

As shown in Table 3, introducing the LIPM and DPPS modules to the baseline model leads to clear improvements in recommendation performance. This confirms the effectiveness of each individual module. The baseline model achieves 0.436 in NDCG@10, 0.556 in Recall@10, and 0.344 in

Precision@10. These results serve as reference points, reflecting the performance without personalized aggregation or privacy protection mechanisms.

When only the Local Interest-guided Personalized Model aggregation module (+LIPM) is added, all three metrics improve. In particular, NDCG@10 increases to 0.462. This indicates that the model captures user ranking preferences more effectively. LIPM guides parameter aggregation based on user interests. This helps the global model better adapt to the preference differences across clients, thereby enhancing personalized representation and recommendation relevance.

Adding only the Differential Privacy-driven Personalized Strategy (+DPPS) also improves performance. The gains are especially noticeable in Recall@10 and Precision@10. This shows that DPPS preserves strong user preference modeling ability while protecting privacy. By dynamically adjusting noise intensity and gradient guidance, the module reduces the negative impact of privacy mechanisms on model performance. It helps maintain a balance between recommendation accuracy and privacy security.

The full model (Ours), which integrates both LIPM and DPPS, achieves the best results across all metrics. NDCG@10 reaches 0.481, Recall@10 reaches 0.595, and Precision@10 increases to 0.374. These results further demonstrate the synergistic effect of the two modules. They confirm that the proposed collaborative optimization strategy not only enhances personalized recommendation quality but also ensures robust user privacy protection.

3) The impact of non-IID degree on recommendation performance in federated learning

This paper also gives the research results on the impact of non-IID degree on recommendation performance in federated learning, as shown in Figure 4.

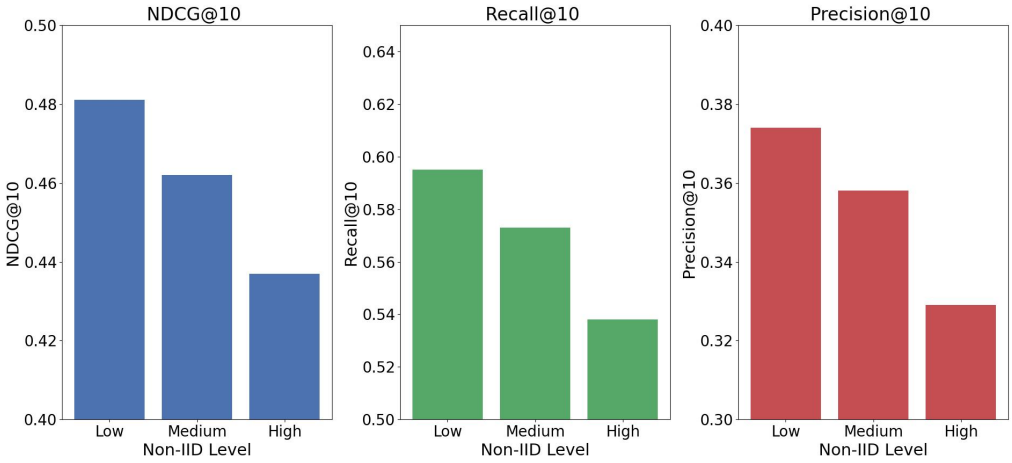

Figure 4. The impact of non-IID degree on recommendation performance in federated learning

As shown in Figure 4, the performance of the federated recommendation system declines noticeably across evaluation metrics as the degree of data non-independence and non-identical distribution (Non-IID) increases. For NDCG@10, the system performs best under the Low Non-IID setting, reaching about 0.48. This value drops to around 0.44 under the High Non-IID condition. This indicates that greater interest divergence among users negatively affects the accuracy of recommendation ranking.

In terms of Recall@10, the recall rate is approximately 0.60 under Low Non-IID, but decreases to around 0.54 under High Non-IID. This suggests that as client data become more diverse, the system finds it harder to capture content that users are truly interested in. The result highlights the global model's limited adaptability when facing inconsistent user preferences.

Precision@10 shows a similar pattern. It decreases from 0.37 under Low Non-IID to about 0.33 under High Non-IID. This drop reflects a reduced presence of relevant items in the recommendation list. Even with a fixed list length, the system's ability to accurately recommend related content is weakened by data heterogeneity, which degrades the overall user experience and precision.

These results demonstrate that Non-IID characteristics of data have a significant impact on recommendation performance in federated learning environments. Without effective personalization mechanisms, the model struggles to maintain stable recommendation quality when exposed to highly heterogeneous data. Therefore, it is essential to design strategies with strong personalization adaptability for federated recommendation systems.

4) The impact of federation aggregation rounds on personalized recommendation effects

This paper also gives the impact of the number of federated aggregation rounds on the personalized recommendation effect, and the experimental results are shown in Figure 5.

Figure 5 illustrates the impact of the number of federated aggregation rounds on personalized recommendation performance. As the number of rounds increases, the model shows steady improvements in NDCG@10, Recall@10, and

Precision@10. This indicates that more communication rounds help the model better integrate personalized information from clients, thereby enhancing the expressiveness and generalization of the global model.

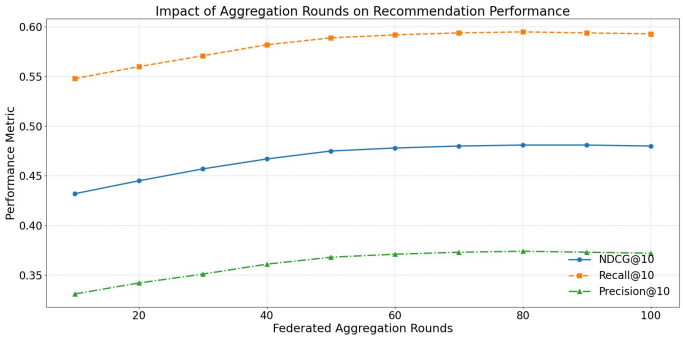

Figure 5. The impact of federation aggregation rounds on personalized recommendation effects

For Recall@10, the curve rises quickly and then saturates. This suggests that the model can significantly improve recall with relatively few rounds, but the performance gain becomes marginal after a certain point. It indicates that federated learning captures user preferences effectively during early aggregation, while further rounds yield limited benefit.

The increase in NDCG@10 is more gradual and consistent. The improvement is stable across iterations. This means that ranking quality continues to improve as the model is trained further. The trend shows that sufficient communication rounds are important for boosting the relevance and accuracy of ranked recommendation lists, especially when aiming for high-quality personalized results.

Precision@10 also shows a slow but steady upward trend. This further confirms that recommendation accuracy improves with more rounds of communication. Overall, the results demonstrate the critical role of aggregation rounds in federated personalized recommendation. Properly tuning the communication frequency helps achieve a balance between efficiency and performance.

5) Analysis of the interference of noise injection mechanism on user preference modeling ability

This paper also presents the experimental results of the interference analysis of the noise injection mechanism on the user preference modeling ability, as shown in Figure 6.

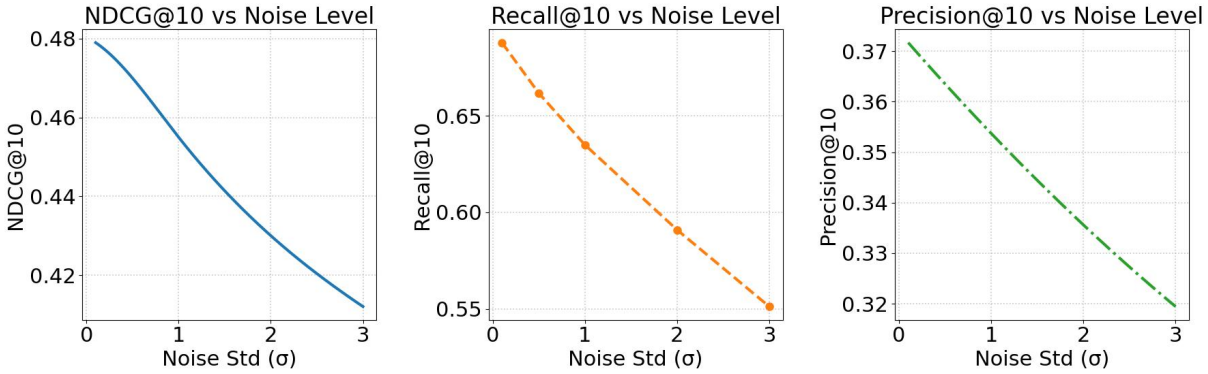

Figure 6. Analysis of the interference of noise injection mechanism on user preference modeling ability

Figure 6 shows the impact of noise injection on the model's ability to capture user preferences. As the standard deviation (σ) of the noise increases, all three main evaluation metrics—NDCG@10, Recall@10, and Precision@10—show a downward trend. This indicates that stronger noise introduces significant interference in personalized recommendation performance. At low noise levels, the model still maintains good ranking and recommendation quality. However, as the noise becomes stronger, performance degradation becomes more pronounced.

The NDCG@10 curve decreases smoothly, indicating that ranking performance is highly sensitive to noise. Since ranking metrics depend on fine-grained modeling of user preferences, Gaussian noise disrupts the model's ability to capture detailed interests. This results in high-quality items being misranked or overlooked, which affects the overall recommendation experience.

Recall@10 drops more sharply. This suggests that under stronger noise, the model fails to effectively retrieve items that users are truly interested in. Noise injection reduces the model's ability to identify highly relevant content, limiting the coverage of the recommendation list. This presents a significant risk for applications where high recall is essential.

The Precision@10 curve also declines gradually, showing that stronger noise directly harms recommendation accuracy. With a fixed recommendation list length, noisy predictions increase the proportion of irrelevant items, which lowers precision. Overall, this experiment reveals that while differential privacy mechanisms help protect user data, they may also degrade model performance. It highlights the need to carefully control noise intensity to balance privacy protection and personalized modeling.

6) Loss function changes with epoch

Finally, this paper presents a graph that illustrates the change in the loss function over training epochs, as shown in Figure 7. The purpose of this visualization is to provide a clear understanding of the model's optimization process during training. By tracking the loss value across iterations, it becomes possible to evaluate how effectively the model minimizes error over time.

This graph serves as an important indicator of training dynamics, including convergence behavior and model stability. Observing the trend of the loss curve helps determine whether the model is learning efficiently and whether it has reached a stable state. It also provides insight into the quality of the training process and the appropriateness of the model configuration

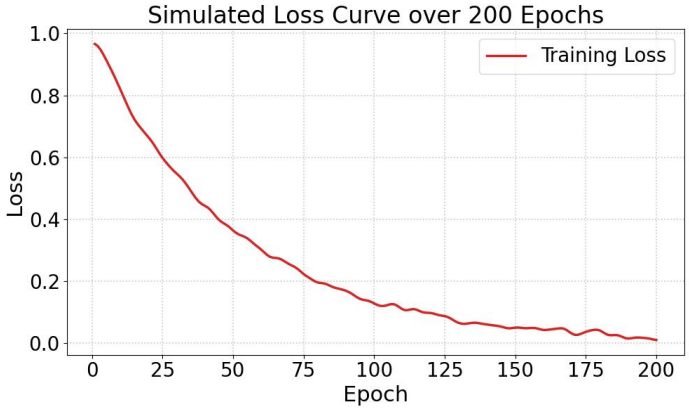

Figure 7. Loss function changes with epoch

As shown in Figure 7, the training loss shows a stable downward trend over 200 epochs. This indicates that the model continuously updates its parameters to minimize error during the iterative optimization process. In the early stages, the loss decreases rapidly. This reflects the model's quick adaptation to global patterns and correction of initial prediction errors, which is a common characteristic in deep learning training.

In the middle and later stages of training, the loss curve gradually flattens. As the loss approaches zero, the model is considered to have reached a good level of convergence. After around 150 epochs, the changes in loss become minimal. This suggests that further learning has limited impact on reducing overall error, marking the beginning of the convergence phase. It demonstrates the training process is both stable and convergent. Moreover, the curve does not show significant oscillations or rebounds. This implies that the training process does not suffer from overfitting or gradient instability. Such behavior typically reflects a well-designed model structure and properly configured optimizer settings. It also indicates that the model has strong potential to maintain good performance on test or real-world data. Overall, the loss curve confirms that the training process is efficient and well-converged.

5. Conclusion

This paper focuses on the conflict between privacy protection and personalized modeling in federated recommendation systems. It proposes a collaborative optimization framework that integrates a local interest-guided personalized aggregation mechanism and a differential privacy-driven update strategy. This enhances its modeling capability in non-independent and identically distributed environments. Experimental results confirm the superiority of the proposed method across multiple evaluation metrics, demonstrating its ability to balance personalization and privacy protection. From a technical perspective, this study addresses limitations of existing methods, such as insufficient personalization or severe performance degradation due to privacy interference. It does so by introducing interest-driven aggregation and privacy-adaptive regulation within a federated learning architecture. The two modules complement each other, enabling the model to maintain stable recommendation performance in heterogeneous data scenarios. The approach avoids the shortcomings of unified modeling, which ignores user diversity, and overcomes information loss caused by noise injection. This structured optimization strategy provides a replicable and scalable theoretical foundation for future federated recommendation system design.

In terms of practical application, the proposed method can be widely applied in privacy-sensitive domains with strong personalization demands, such as healthcare, financial risk control, and smart retail. By delivering high-quality recommendations without sharing raw data, related systems can offer more intelligent and trustworthy services while remaining compliant with data regulations. Furthermore, the framework's generality offers insights for other multi-party collaborative modeling tasks and shows potential for cross-domain adoption. Future research may further explore robust modeling under complex conditions such as heterogeneous devices, dynamic communication frequencies, and cross-platform collaboration. It may also consider emerging challenges in federated recommendation, including fairness and interpretability. As privacy-preserving technologies advance and real-world deployments expand, building efficient, secure, and personalized recommendation systems will become a critical direction for intelligent systems. This study represents a key step toward that goal.

References

[1] Yang, Liu, et al. "Federated recommendation systems." Federated Learning: Privacy and Incentive. Cham: Springer International Publishing, 2020. 225-239.

[2] Li, Zhiwei, et al. "Navigating the future of federated recommendation systems with foundation models." arXiv preprint arXiv:2406.00004 (2024).

[3] Yu, Enqi, et al. "A federated recommendation algorithm based on user clustering and meta-learning." Applied Soft Computing 158 (2024): 111483.

[4] Qu, Liang, et al. "Towards personalized privacy: User-governed data contribution for federated recommendation." Proceedings of the ACM Web Conference 2024. 2024.

[5] Liu, S., Xu, S., Yu, W., Fu, Z., Zhang, Y., & Marian, A. (2021, July). FedCT: Federated collaborative transfer for recommendation. In Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval (pp. 716-725).

[6] Zhang, Honglei, et al. "Transfr: Transferable federated recommendation with pre-trained language models." arXiv preprint arXiv:2402.01124 (2024).

[7] Sun, Zehua, et al. "A survey on federated recommendation systems." IEEE Transactions on Neural Networks and Learning Systems 36.1 (2024): 6-20.

[8] Wu, C., Wu, F., Cao, Y., Huang, Y., & Xie, X. (2021). Fedgnn: Federated graph neural network for privacy-preserving recommendation. arXiv preprint arXiv:2102.04925.

[9] Zhang, Honglei, et al. "Privfr: Privacy-enhanced federated recommendation with shared hash embedding." IEEE Transactions on Neural Networks and Learning Systems (2024).

[10] Li, Zhitao, et al. "Decentralized federated recommendation with privacy-aware structured client-level graph." ACM Transactions on Intelligent Systems and Technology 15.4 (2024): 1-23.

[11] Pei, Jiaming, et al. "A review of federated learning methods in heterogeneous scenarios." IEEE Transactions on Consumer Electronics (2024).

[12] Chen, Jingxue, et al. "When federated learning meets privacy-preserving computation." ACM Computing Surveys 56.12 (2024): 1-36.

[13] Ali, Saqib, Qianmu Li, and Abdullah Yousafzai. "Blockchain and federated learning-based intrusion detection approaches for edge-enabled industrial IoT networks: A survey." Ad Hoc Networks 152 (2024): 103320.

[14] Yazdinejad, Abbas, et al. "A robust privacy-preserving federated learning model against model poisoning attacks." IEEE Transactions on Information Forensics and Security (2024).

[15] Ficco, Massimo, et al. "Federated learning for IoT devices: Enhancing TinyML with on-board training." Information Fusion 104 (2024): 102189.

[16] Qi, Pian, et al. "Model aggregation techniques in federated learning: A comprehensive survey." Future Generation Computer Systems 150 (2024): 272-293.

[17] Wu, Siyu, et al. "A comprehensive exploration of personalized learning in smart education: From student modeling to personalized recommendations." arXiv preprint arXiv:2402.01666 (2024).

[18] Wagner, K., Merceron, A., Sauer, P., & Pinkwart, N. (2022). Personalized and explainable course recommendations for students at risk of dropping out. In Proceedings of the 15th International Conference on Educational Data Mining (p. 657).

[19] Luo, S., Xiao, Y., & Song, L. (2022, October). Personalized federated recommendation via joint representation learning, user clustering, and model adaptation. In Proceedings of the 31st ACM international conference on information & knowledge management (pp. 4289-4293).

[20] Gm, D., Goudar, R. H., Kulkarni, A. A., Rathod, V. N., & Hukkeri, G. S. (2024). A digital recommendation system for personalized learning to enhance online education: A review. IEEE Access, 12, 34019-34041.

[21] Huang, Shuaishuai, et al. "Deep adaptive interest network: personalized recommendation with context-aware learning." arXiv preprint arXiv:2409.02425 (2024).

[22] Lin, Guanyu, et al. "Fedrec: Federated recommendation with explicit feedback." IEEE Intelligent Systems 36.5 (2020): 21-30.

[23] Chai, Di, et al. "Secure federated matrix factorization." IEEE Intelligent Systems 36.5 (2020): 11-20.

[24] Jiang, Xueyong, et al. "FedNCF: federated neural collaborative filtering for privacy-preserving recommender system." 2022 International joint conference on neural networks (IJCNN). IEEE, 2022.

[25] Cheng, Anda, et al. "Differentially private federated learning with local regularization and sparsification." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2022.

[26] T Dinh, Canh, Nguyen Tran, and Josh Nguyen. "Personalized federated learning with moreau envelopes." Advances in neural information processing systems 33 (2020): 21394-21405.