PDF Download
3742791.pdf
21 December 2025
Total Citations: 0
Total Downloads: 694

Published: 23 July 2025
Online AM: 06 June 2025
Accepted: 05 May 2025
Revised: 07 April 2025
Received: 13 November 2024

Citation in BibTeX format

RESEARCH-ARTICLE

A Federated Graph Neural Network with Differential Privacy for Cross-domain Recommender Systems

PHAM MINH DO, University of Technology Sydney, Sydney, NSW, Australia

JIE LU, University of Technology Sydney, Sydney, NSW, Australia

QIAN ZHANG, University of Technology Sydney, Sydney, NSW, Australia

GUANGQUAN ZHANG, University of Technology Sydney, Sydney, NSW, Australia

ACM Transactions on Intelligent Systems and Technology, Volume 16, Issue 4 (August 2025)
https://doi.org/10.1145/3742791
EISSN: 2157-6912

A Federated Graph Neural Network with Differential Privacy for Cross-domain Recommender Systems

PHAM MINH THU DO, JIE LU, QIAN ZHANG, and GUANGQUAN ZHANG, Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, Australia

Cross-domain recommender systems, which are designed to address issues with data sparsity, tend to suffer notable challenges with safeguarding user privacy. While existing cross-domain recommendation methods incorporate privacy mechanisms, they often fall short in practice, offering only one-sided benefits and limited privacy safeguards. In this study, we propose a novel privacy-preserving cross-domain recommender system that combines federated transfer learning with differential privacy to facilitate cross-domain knowledge transfer while ensuring strong privacy protection. First, we leverage federated transfer learning, treating each domain as an independent client to protect privacy for business partners by preventing the exchange of raw data. Second, we use a graph neural network (GNN) as the encoder to learn the user and item representations. We also design a consistency loss function that maintains the invariance between local and global user representations while preventing representation collapse. Third, we introduce a privacy mechanism that applies differential privacy to the output of each aggregation layer in the GNN—the aim being to protect transferred user representations while balancing privacy with accuracy. Finally, our transfer mechanism operates without user-identifying information, establishing connections between domains by detecting latent overlapping users and subsequently performing personalized preference aggregation. This allows for efficient knowledge transfer across domains. Experiments on real-world datasets show that our approach significantly enhances recommendation accuracy while offering robust privacy protection, outperforming state-of-the-art baselines.

CCS Concepts: • **Information systems** → **Recommender systems**; • **Security and privacy**;

Additional Key Words and Phrases: cross-domain recommendation, privacy-preserving, graph neural network, federated learning, differential privacy

ACM Reference format:
Pham Minh Thu Do, Jie Lu, Qian Zhang, and Guangquan Zhang. 2025. A Federated Graph Neural Network with Differential Privacy for Cross-domain Recommender Systems. *ACM Trans. Intell. Syst. Technol.* 16, 4, Article 88 (July 2025), 25 pages.
https://doi.org/10.1145/3742791

Funding support for this article was provided by the Australian Research Council (ARC) (FL190100149, DP220102635).
Authors' Contact Information: Pham Minh Thu Do, Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, Australia; e-mail: thu.p.do@student.uts.edu.au; Jie Lu (corresponding author), Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, Australia; e-mail: jie.lu@uts.edu.au; Qian Zhang, Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, Australia; e-mail: Qian.Zhang-1@uts.edu.au; Guangquan Zhang, Australian Artificial Intelligence Institute, University of Technology Sydney, Sydney, Australia; e-mail: guangquan.zhang@uts.edu.au.

ACM 2157-6912/2025/7-ART88
https://doi.org/10.1145/3742791

1 Introduction

Collaborative filtering [16, 24, 27, 34] is a widely used technique for building recommender systems based on user preferences. However, its performance significantly deteriorates in the presence of data sparsity. **Cross-domain recommender systems (CDRSs)**, which enable knowledge transfer across multiple domains, have emerged as an effective solution to alleviate this issue. Privacy regulations impose a major challenge in developing practical CDRSs. For example, in digital advertising, companies aim to exchange insights on user preferences to enhance targeted advertising while adhering to strict privacy policies. In online education, multiple institutions may wish to collaborate on personalized course recommendations without exposing student identities. However, traditional CDRSs facilitate cross-domain knowledge transfer either by sharing raw user behavior data [13, 20, 44] or by exchanging extracted user representations [2, 21], neither of which fully complies with privacy requirements. In practice, data from different domains are often owned by separate entities, and privacy regulations strictly prohibit direct data sharing [47]. Moreover, transferring high-quality user representations can still pose significant privacy risks, leading to potential information leakage [3, 39].

As a result, privacy-preserving CDRSs have gained increasing attention. Several recent methods attempt to address privacy concerns through various techniques. PriCDR [4] employs **differential privacy (DP)** to safeguard the source rating matrix, while PPGenCDR [18] leverages a conditional **generative adversarial network (GAN)** with DP to protect user preferences. FedCDR [26] utilizes federated learning with personalized strategies to improve recommendations in the target domain. However, these methods primarily focus on unidirectional knowledge transfer, which can lead to negative transfer when the direction is reversed, discouraging source-domain data owners from participating in real-world collaborations. Additionally, FedCDR treats edge devices as clients that store local user data, which may not fully address privacy concerns in inter-company collaborations, where businesses require robust privacy protection. P2FCDR [5] partially addresses these limitations by integrating federated learning with DP to protect user representations. However, its reliance on deep matrix factorization as an encoder may restrict its capability to learn high-quality user representations. PPCDR [37] introduces a graph-based architecture with DP to enhance protection before exchanging user representations across domains.

Despite the progress of existing methods, several limitations hinder their practical applicability. One major challenge is their reliance on user-identifying information to detect overlapping users. This dependency severely limits practical feasibility in privacy-sensitive contexts, particularly when a shared user identifier is absent. For example, in healthcare recommender systems, patient identities must remain strictly confidential, making direct user matching impossible. Another critical challenge stems from the lack of differentiation between local and global user representations. While incorporating both representations improves user preference modeling, failing to properly distinguish them can lead to negative transfer, where domain-specific information is erroneously shared across domains. Thus, three fundamental challenges in building privacy-preserving CDRS are: (1) Designing a privacy-preserving framework that benefits all participants while effectively balancing privacy and accuracy, (2) learning high-quality user representations while mitigating the risk of negative transfer, and (3) developing a general and stable knowledge transfer mechanism that is adaptable to multiple scenarios.

In this study, we propose a novel approach, Federated Graph neural network with Differential privacy for Cross-Domain Recommender systems, called FGD-CDR for short. Our FGD-CDR method combines federated transfer learning with DP to facilitate cross-domain knowledge transfer without compromising user privacy. First, we implement federated transfer learning, treating each domain as a client to ensure that no original data are shared across domains. Second, within each domain, we extract both local and global user representations, along with item representations, from the interaction graph. A specially designed consistency loss encourages consistent information capture between local and global user representations while ensuring they remain distinct and discriminative. Next, we introduce a privacy mechanism in which DP is applied after each aggregation step within the GNN to protect the global user representations shared between domains. Lastly, a transfer mechanism establishes connections and aggregates preferences among domains, guiding the knowledge transfer process. The **connection establishment (CE)** identifies latent overlapping users, creating implicit connections between domains, while the **personalized preference aggregation (PPA)** ensures domain-specific adaptation, generating enhanced global user representations.

In summary, this article makes the following contributions:

- We introduce a novel CDRS that incorporates robust privacy protection designed for real-world applications. Our method supports bidirectional knowledge transfer and treats each domain as an independent client, ensuring that all data remains locally stored.
- We introduce a consistency loss that enhances alignment between local and global user representations while maintaining their distinctiveness. This improves representation learning and reduces the risk of transferring irrelevant knowledge.
- We propose a privacy mechanism that applies Laplace noise to the output of the aggregation function, safeguarding user representations before transfer. This approach preserves the essential structure of the input graph, provides robust privacy protection, and better balances privacy with accuracy.
- We present a novel transfer mechanism that identifies latent overlaps in users to establish inter-domain connections and performs personalized aggregation for domain-specific adaptation. Hence, our approach does not rely on user-identifying information, making it highly suitable for privacy-sensitive environments. Additionally, this transfer mechanism is adaptable to various cross-domain scenarios, including both overlapping and non-overlapping user settings.

The remainder of this article is structured as follows: Section 2 offers a comprehensive overview of the relevant works. Section 3 introduces essential background concepts and outlines the problem statement. In Section 4, we detail our proposed methodology, focusing on the overall framework and its individual components. Section 5 presents and analyzes the experimental results, including an ablation study to assess the contributions of various components. Finally, Section 6 summarizes our work and discusses potential directions for future research.

2 Related Works

This section provides an overview of current CDRSs, with a focus on approaches that prioritize privacy preservation.

2.1 CDRSs

CDRSs have emerged as a powerful solution to the challenge of data sparsity by enabling knowledge exchange across different domains. The transferred knowledge can either be user interaction data [6, 17, 43, 44] or extracted user representations [2, 21, 23, 46]. For example, PPGN [44] constructs a cross-domain interaction graph to model interactions across domains. HeroGraph [6] employs recurrent attention combined with graph convolution operations to embed the cross-domain interaction graph into a latent space. The representations learned from the shared structure are embedded in a common coordinate system, enabling seamless knowledge transfer across domains. DACDR [43] enhances representation learning by capturing both local and global representations from local interaction matrices and cross-domain interaction matrices, respectively. Its adversarial network, with a domain discriminator, prevents domain-specific knowledge from being included in the global representations, thus avoiding negative transfer.

Recent studies have concentrated on transferring user representations extracted from each domain. EMCDR [23] maps source user representations to the target domain using a non-linear mapping function. GA-DTCDR [46] facilitates knowledge transfer by integrating the source and target representations of overlapping users through an attention mechanism. Bi-TGCF [21] incorporates **graph neural networks (GNNs)** for representation learning, enabling effective knowledge transfer between graph convolutional networks. DisenCDR [2] disentangles cross-domain and domain-specific information, effectively exploiting common patterns across domains. Although these works achieve impressive performance, they face a significant challenge concerning privacy protection. The absence of robust privacy-preserving techniques hinders their ability to adapt to real-world scenarios, particularly in sensitive domains where the confidentiality of user data is critical [38].

2.2 Privacy-preserving CDRSs

The development of privacy-preserving techniques for CDRSs has become a critical focus, especially in the context of knowledge transfer. One of the pioneering works addressing privacy concerns in CDRS is NATR [9], which transfers item-side information instead of sensitive user data to protect privacy. DP [7] has been widely adopted to protect data prior to transfer. For example, PriCDR [4] employs a DP algorithm to secure the source rating matrix before transferring it to the target domain. PPGenCDR [18] introduces a privacy-preserving generator using a conditional GAN model, where the generator produces synthetic user preferences with DP, and the discriminator distinguishes between real and synthetic preferences.

Federated learning [14] enables collaborative model training among multiple participants while preserving the privacy of their raw data. This setup allows each participant to train a local model using their own data and transmits only the local model updates to a central server. The server then updates the global model by aggregating these local parameters, which are subsequently distributed back to all participants. This iterative process continues until the global model converges, facilitating collective learning without compromising individual data privacy. Based on data distribution among participants, federated learning is categorized into three types: horizontal [19], vertical [22], and federated transfer learning [32]. Federated learning is widely adopted as a core element in designing privacy-preserving CDRSs. For example, FedCDR [26] applies federated learning technique, with personalized update and aggregation strategies tailored to each client and server. To ensure privacy, only non-personal parameters, such as item representations and transfer modules, are transmitted to the server. Recent research has focused on federated transfer learning, which combines the principles of transfer learning and federated learning. For example, P2FCDR [5] is a dual-target CDR that applies DP to user representations before transferring them across domains using an orthogonal mapping function. Similarly, PPCDR [37] introduces a graph transfer module that merges global and local user representations during the private update process, applying DP to protect shared user representations prior to transfer.

3 Problem Definition and Preliminaries

This section delineates the theoretical preliminaries, followed by the research problem addressed in this article.

3.1 Theoretical Preliminaries

3.1.1 GNNs. GNNs [40] aim to learn node representations by integrating initial node features with a graph structure (edges). A standard K-layer GNN comprises K graph convolution layers.

In layer k, the nodes receive representations from layer $k-1$, denoted by X^{k-1}, and undergo an aggregation step followed by an update step to produce new node representations X^k:

$$X^k = UPD(AGG(\mathsf{A}, X^{k-1}); \Theta^k). \tag{1}$$

The function $AGG(\cdot)$ serves as the aggregation function, which can take forms such as mean, sum, or max pooling. $UPD(\cdot)$ is the update function, typically a neural network-based transformation parameterized by Θ^k.

3.1.2 DP. DP [8] is a fundamental principle in data privacy that ensures any computation's output based on a dataset does not allow an attacker to infer information about any individual in that dataset.

Definition 1 (ϵ-DP [7]). An algorithm $\mathcal{A}$ satisfies ϵ-differential privacy (ϵ-DP), where $\epsilon \geq 0$, if and only if, for any set of outputs $O \subseteq \text{Range}(\mathcal{A})$ and any neighboring datasets D and D' that differ by one element, it holds that

$$\Pr[\mathcal{A}(D) \in O] \leq \exp(\epsilon) \Pr[\mathcal{A}(D') \in O],$$

where exp(.) denotes the exponential function, and $\text{Range}(\mathcal{A})$ is the set of all possible outputs of the algorithm $\mathcal{A}$. The parameter ϵ, known as the privacy budget, controls the strength of privacy: a smaller ϵ provides stronger privacy protection, while a larger ϵ weakens it.

Definition 2 (Sensitivity [7]). The sensitivity of the function $f : D \longrightarrow \mathbb{R}$ is defined as:

$$\Delta_f = \max_{D,D'} \|f(D) - f(D')\|_2.$$

Sensitivity quantifies the maximum possible change in the output of f when applied to two neighboring datasets D and D', differing by a single element. In this case, it refers to the maximum $L2$ distance between the function outputs for these two datasets.

A Laplace mechanism that adds noise to the function's output based on the function's sensitivity is a widely used method for achieving DP.

Definition 3 (Laplace Mechanism [7]). Given the function $f : D \longrightarrow \mathbb{R}^d$, the following mechanism $\mathcal{A}$ satisfies ϵ-DP:

$$\mathcal{A}(D) = f(D) + \text{Lap}\left(\frac{\Delta_f}{\epsilon}\right)^d,$$

where $\text{Lap}\left(\frac{\Delta_f}{\epsilon}\right)^d$ denotes a random variable drawn from a d-dimensional Laplace distribution with a zero mean and a scale of $\frac{\Delta_f}{\epsilon}$.

Graph datasets differ from traditional tabular datasets because they also contain data about the connections between data records. Therefore, defining edge-level adjacency in graphs is essential.

Definition 4 (Edge-level Adjacent Graphs [10]). Two graphs $\mathcal{G}$ and $\mathcal{G}'$ are considered edge-level adjacent if one can be formed from the other by adding or removing a single edge. Thus, $\mathcal{G}$ and $\mathcal{G}'$ differ by at most one edge.

Accordingly, the definition of edge-level DP is as follows:

Definition 5 (Edge-level ϵ-DP [33]). An algorithm $\mathcal{A}$ satisfies edge-level ϵ-DP, where $\epsilon \geq 0$, if and only if, for any set of outputs $O \subseteq \text{Range}(\mathcal{A})$ and any neighboring graphs $\mathcal{G}$ and $\mathcal{G}'$ that differ by

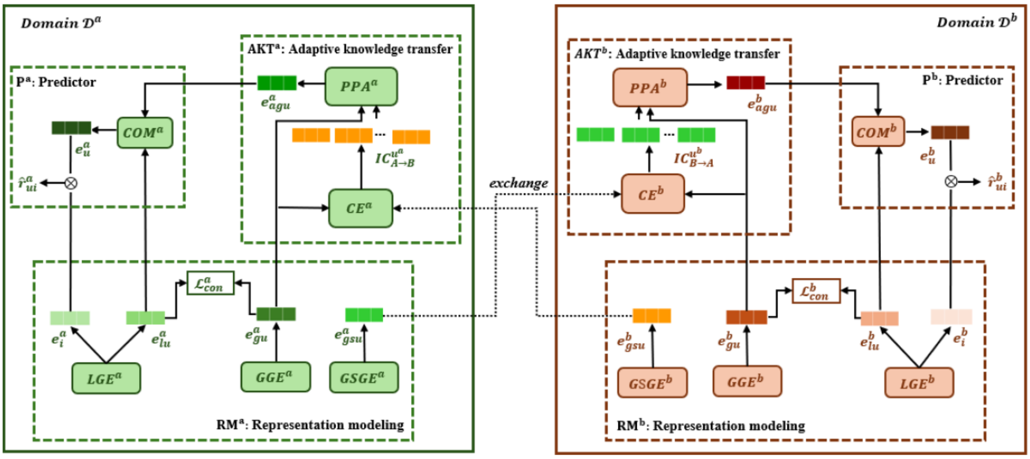

Fig. 1. The network architecture of FGD-CDR.

one edge, the level of privacy satisfies

$$\Pr[\mathcal{A}(\mathcal{G}) \in O] \leq \exp(\epsilon) \Pr[\mathcal{A}(\mathcal{G}') \in O].$$

This definition ensures that the adjacency relationship between datasets accounts for the structural changes in the graph, specifically focusing on the connections between nodes.

3.2 Problem Definition

We address the cross-domain recommendation problem by considering two domains, denoted as $\mathcal{D}^a$ and $\mathcal{D}^b$. Each domain includes sets of users and items, denoted as $A = \{\mathcal{U}^a, \mathcal{I}^a\}$ for $\mathcal{D}^a$ and $B = \{\mathcal{U}^b, \mathcal{I}^b\}$ for $\mathcal{D}^b$. We assume the two domains share the same set of users, represented by $\mathcal{U}$. The number of users, items in domain $\mathcal{D}^a$, and items in domain $\mathcal{D}^b$ are n_u, n_i^a and n_i^b, respectively. The user-item interaction graphs in $\mathcal{D}^a$ and $\mathcal{D}^b$ are denoted as $\mathcal{G}^a = (\mathcal{U}, \mathcal{I}^a, \mathcal{R}^a)$ and $\mathcal{G}^b = (\mathcal{U}, \mathcal{I}^b, \mathcal{R}^b)$, respectively. Here, $\mathcal{R}^a \in \mathbb{R}^{n_u \times n_i^a}$ and $\mathcal{R}^b \in \mathbb{R}^{n_u \times n_i^b}$ are the interaction matrices, where the entry r_{jv} is 1 if user u_j interacts with item i_v.

Our primary objective is to enhance top-N recommendation performance in participating domains while preserving user privacy under the following constraints:

- *Rule 1*: The original interaction data are not shared between domains.
- *Rule 2*: The user representations learned from the interaction graph in each domain must be protected before exchange.

Due to this privacy-preserving setting, even if user/item overlap exists across domains, identifying information remains difficult to acquire. Furthermore, since interactions in each domain are modeled as an interaction graph, the privacy mechanism in each domain must satisfy edge-level DP.

4 Methodology

4.1 Overview of Our FGD-CDR

Building on recent advancements in federated learning [15, 36] and DP [41], we propose a privacy-preserving framework for cross-domain recommendation. The architecture of our proposed method, FGD-CDR, is shown in Figure 1, with key notations summarized in Table 1.

Table 1. Frequently Used Notations

Notations	Description
LGE^a, LGE^b	Local graph encoders
GGE^a, GGE^b	Global graph encoders
$GSGE^a, GSGE^b$	Global-shared graph encoders
CE^a, CE^b	Connection establishment
PPA^a, PPB^b	Personalized preference aggregation
e^a_{gu}, e^b_{gu}	Global user representations
e^a_{lu}, e^b_{lu}	Local user representations
e^a_{gsu}, e^b_{gsu}	Global-shared user representations
e^a_{agu}, e^b_{agu}	Adapted global user representations
e^a_u, e^b_u	Final user representations
e^a_i, e^b_i	Item representations
$\hat{r}^a_{ui}, \hat{r}^b_{uj}$	Ranking scores
$\mathcal{L}^a_{con}, \mathcal{L}^b_{con}$	Consistency losses

Our proposed framework is composed of the following main components:

- *Representation Modeling (RM)*: In each domain, item representations, local user representations, and global user representations are extracted from the local interaction graph. To ensure alignment between the local and global user representations while preserving their individual characteristics, we introduce a consistency loss. Furthermore, DP is applied during the aggregation process, which generates the global-shared user representations. These protected representations are then exchanged between domains to facilitate knowledge transfer.
- *Adaptive Knowledge Transfer (AKT)*: Each domain receives global-shared user representations from other domains. To ensure privacy, these representations are exchanged without any user-identifying information. We introduce a knowledge transfer mechanism that incorporates CE to identify latent overlapping users and establish connections between domains, followed by PPA to generate adapted global user representations.
- *Predictor (P)*: The adapted global user representations are combined with local user representations to create the final user representations. A ranking score is then calculated as the inner product of the final user representation and the item representation.

Our FGD-CDR framework combines DP with federated transfer learning to strengthen privacy protection in CDRSs. In contrast to traditional methods that focus on individual user data privacy, our framework emphasizes protecting the privacy of business partners by treating each domain as an independent client, making it particularly suited for real-world applications. First, each domain retains its data locally without directly sharing it, significantly reducing the risk of privacy leaks [25]. Second, directly transferring high-quality user representations could expose sensitive user data [3]. To counter this risk, we employ DP by introducing controlled Laplace noise to user representations prior to any exchange. Although a smaller privacy budget (ϵ) increases privacy protection, it may slightly impact the quality of the transferred representations, necessitating a careful balance between privacy and performance. Lastly, to further enhance privacy, only global-shared user representations, devoid of any identifying information, are exchanged, preventing the detection of overlapping users across domains.

4.2 RM

This section introduces the RM approach, which learns item representations, along with local, global, and global-shared user representations, from the local interaction graph. It incorporates a consistency loss and a privacy mechanism to ensure both effective learning and robust protection.

4.2.1 Local Graph Encoder (LGE) and Global Graph Encoder (GGE). Building on the effectiveness of GNNs in representation learning [11, 42], we design an LGE and a GGE integrated with a consistency loss. This design enables capturing both local and global user representations from the user-item interaction graph within each domain. Specifically, local user representations focus on domain-specific preferences and characteristics tied to each domain, while global user representations identify more generalizable preferences that span across multiple domains. In real-world scenarios, users frequently exhibit overlapping preferences across domains while also retaining distinct interests unique to each domain. Our framework, by jointly modeling both local and global representations, achieves a balance between capturing cross-domain patterns and preserving the unique aspects of user behavior within each domain. To achieve these representations, we leverage the propagation and aggregation mechanisms of LightGCN [11], an efficient model for capturing user-item interactions.

4.2.2 Consistency Loss. Since local and global user representations relate to the same user, they should closely align but still retain their distinct information. To enforce this, we introduce a consistency loss that ensures local and global representations are aligned, but not identical, to prevent feature collapse. The consistency loss is formulated as:

$$\mathcal{L}_{con} = \frac{1}{n_u} \sum_{i=1}^{n_u} ||e_{gu_i} - e_{lu i}||^2 + \gamma \cdot log \frac{1}{n_u(n_u - 1)} \sum_{i=1}^{n_u} \sum_{j=1, j \neq i}^{n_u} e^{-t||e_{gu_i} - e_{lu j}||^2}, \tag{2}$$

where $e^{-t||e_{gu_i} - e_{lu j}||^2}$ represents the exponential of the negative squared Euclidean distance between global and local user representations, scaled by a factor t. In the above equation, the first term promotes alignment between the global user representation e_{gu} and the local user representation e_{lu}, ensuring consistency. The second term encourages a uniform distribution of the representations, preventing feature collapse and ensuring that local and global user representations remain distinct and discriminable.

This consistency loss guarantees that, while local and global user representations share common information, they retain their unique attributes. Maintaining this distinction is critical to avoid negative transfer, where domain-specific information from one domain negatively impacts recommendations in another [43]. By distinguishing between local and global user representations, domain-specific details are confined to their respective domains, reducing the risk of transferring irrelevant knowledge across domains.

4.2.3 Global-shared Graph Encoder (GSGE). Global user representations encapsulate user preferences that generalize across domains, making them ideal for transfer in cross-domain systems. However, transferring such high-quality representations introduces a significant risk of information leakage, as they may reveal the sensitive historical interactions of the users [3, 39]. Therefore, it is crucial to apply privacy-preserving techniques to global user representations before sharing them across domains.

In edge-level DP, perturbing the edges of an interaction graph can protect the privacy of relationships between users and items [28]. However, altering the graph structure at this level risks distorting the underlying relationships, potentially affecting the interpretability and accuracy of user-item interactions. Moreover, although adding noise directly to the output layer of a model can provide strong privacy guarantees, this approach may overwhelm the learned representations, diminishing their utility and degrading the quality of recommendations. This tension highlights the need for a privacy mechanism that achieves edge-level DP without compromising the meaningful signals embedded in the representations. In other words, we need to design DP -based algorithms that account for DP provenance [12], focusing on protecting user privacy while ensuring that the representations retain valuable knowledge.

To balance privacy and utility, following [33], we introduce a method that applies Laplace noise directly to the output of the model's aggregation function. By perturbing the aggregation step rather than the raw graph data, we ensure that the contributions of the individual edges to the representation are masked, but the overall structure and key patterns remain intact. This approach mitigates the risk of distorting the graph while still providing a robust privacy guarantee.

In domain $\mathcal{D}^a$, the input layer starts with the representation matrix $(X^a)^0$, where each row is normalized to an $L2$-norm of 1:

$$(\check{X}^a)^0 = (X^a)^0 / \|(X^a)^0\|_2. \tag{3}$$

— *Aggregation*: The non-private representation matrix at layer k is computed using the previous private representation matrix $(\check{X}^a)^{k-1}$ and the adjacency matrix A^a.

$$(X^a)^k = AGG((\check{X}^a)^{k-1}, \mathsf{A}^a). \tag{4}$$

— *Perturbation*: Next, we perturb the non-private representation matrix $(X^a)^k$ by adding Laplace noise with a sensitivity of 1 [33] and a privacy budget of ϵ to every row:

$$(\tilde{X}^a)^k = (X^a)^k + Lap(0, 1/\epsilon) \tag{5}$$

where $Lap(0, 1/\epsilon)$ represents noise drawn from a Laplace distribution with mean of 0 and a scale of $1/\epsilon$. This mechanism satisfies ϵ-DP and protects the privacy of individual edges by obscuring the presence or absence of specific edges in the output representations.

— *Normalization*: Finally, we row-normalize the perturbed representation $(\tilde{X}^a)^k$ to obtain $(\check{X}^a)^k$, ensuring each row retains an $L2$-norm of 1, as per Equation (3).

Applying a privacy mechanism after each aggregation step offers several advantages. First, it mitigates the increasing interdependence between node representations, effectively concealing the influence of individual edges in the graph, thereby offering stronger privacy guarantees. Second, the recursive aggregation acts as a smoothing operation, accumulating useful information while averaging out the noise introduced in earlier layers. This method preserves true signals, ensuring that relevant information is retained while minimizing the disruptive effects of noise. The user representation extracted through this process is termed global-shared user representations, which is privacy-protected and suitable for exchange across domains. Algorithm 1 details the steps for generating global user representations in domain $\mathcal{D}^a$. An analogous process is employed in domain $\mathcal{D}^b$.

In our framework, both the GGE and the GSGE utilize the same input user representations. To preserve privacy, we apply Laplace noise to the output of the aggregation function in the GSGE when generating global-shared user representations. This ensures that these representations are protected by DP and are suitable for exchange among domains, thereby facilitating collaborative learning without exposing sensitive information. In contrast, the GGE's outputs remain unchanged and are confined within each domain. The consistency loss promotes alignment between local and global user representations without explicitly sharing values, minimizing the risk of privacy leakage to other domains.

4.3 AKT

To ensure privacy-preserving knowledge transfer in our framework, the global-shared user representations exchanged between domains contain no identifiable user information. After receiving

Algorithm 1: Perturbing the Aggregation Outputs in GSGE at Domain $\mathcal{D}^a$

```
Input  : $(X^a)^0$: input representation matrix
         $\mathsf{A}^a$: adjacency matrix
         $\epsilon$ : privacy budget
         $K$: number of layers
Output : Protected representations $output = \sum(\check{X}^a)^k$ for $k \in K$
Normalization: $(\check{X}^a)^0 = (X^a)^0/\|(X^a)^0\|_2$;
Initialize: $ouput = 0$;
for $k \in K$ do
    Aggregation: $(X^a)^k = AGG((\check{X}^a)^{k-1}, \mathsf{A}^a)$;
    Perturbation: $(\tilde{X}^a)^k = (X^a)^k + Lap(0, 1/\epsilon)$;
    Normalization: $(\check{X}^a)^k = (\tilde{X}^a)^k/\|(\tilde{X}^a)^k\|_2$;
    Accumulate: $output = output + (\check{X}^a)^k$;
end
return output
```

the global-shared representations from other domains, the next critical step is to perform domain adaptation. This adaptation allows each domain to integrate the transferred representations effectively while maintaining privacy.

CE: The first task in this process is to establish connections by identifying latent overlapping users—users who display similar preferences across domains. These users are identified based on the similarity between the global user representations in the current domain and the global-shared user representations from other domains. The key assumption is that generalizable preferences remain stable across domains, even though specific behaviors might vary. For example, the similarity score between user u_i^a in domain $\mathcal{D}^a$ and an unknown user u_j^b in domain $\mathcal{D}^b$ is measured using a Gaussian kernel, as follows:

$$f(u_i^a, u_j^b) = \exp\left(-\frac{||e_{gu_i}^a - e_{gsu_j}^b||^2}{\sigma^2}\right), \tag{6}$$

where $e_{gu_i}^a$ represents the global user representation of user u_i^a in domain $\mathcal{D}^a$, and $e_{gsu_j}^b$ is the global-shared user representation of user u_j^b in domain $\mathcal{D}^b$. The kernel σ^2 controls the sensitivity to differences between the representations.

For given user u_i^a in domain $\mathcal{D}^a$, the k most similar users from domain $\mathcal{D}^b$ are selected based similarity scores.

$$\mathcal{IC}_{\mathcal{A}\to\mathcal{B}}^{u_i^a} = \{u^b \mid f(u_i^a, u^b) \in \textbf{top-}k(\mathcal{F}^{u_i^a})\}, \tag{7}$$

where $\mathcal{F}^{u_i^a} = \{f(u_i^a, u^b) | u^b \in \mathcal{U}\}$ represents the set of similarity scores between user u_i^a and all users in domain $\mathcal{D}^b$. A similar process occurs for users in domain $\mathcal{D}^b$.

PPA: Once the connections are established, we perform PPA in each domain to produce an adapted global user representation. For example, for given user u_i^a in domain $\mathcal{D}^a$, the global user representation $e_{gu_i}^a$ and the top-k most similar global-shared user representations $\mathcal{IC}_{\mathcal{A}\to\mathcal{B}}^{u_i^a}$ from domain $\mathcal{D}^b$ are utilized to generate the adapted global user representation $e_{agu_i}^a$ as follows:

$$e_{agu_i}^a = (1-\alpha)e_{gu_i}^a + \alpha \sum_{u_j^b \in \mathcal{IC}_{\mathcal{A}\to\mathcal{B}}^{u_i^a}} \beta_j e_{gsu_j}^b, \tag{8}$$

where α is the transfer ratio and β_j is the attention coefficient that signifies the contribution of knowledge in $e_{gsu_j}^b$ to domain $\mathcal{D}^a$, calculated using an attention mechanism [46] with a learnable transformation matrix W:

$$\beta_j = \text{softmax}(\text{ReLU}(\text{concat}(e^a_{gu_i}, e^b_{gsu_j}) \cdot W)). \tag{9}$$

Our transfer mechanism establishes robust mappings between domains based on implicit connections without relying on user-identifying information, making it particularly well-suited for privacy-sensitive applications. This mechanism is highly adaptable, accommodating diverse cross-domain configurations, including both overlapping and non-overlapping user scenarios. Additionally, establishing implicit connections through latent overlapping user groups helps to reduce noise from users interested only in items within their current domain.

4.4 Predictor

Local user representations capture preferences that are specific to a particular domain, providing insights into individual user behavior within that context. In contrast, adapted global user representations incorporate knowledge from other domains, reflecting common preferences shared across multiple domains. By combining these two types of representations, we construct a comprehensive view of user preferences that balances unique domain-specific interests with shared, cross-domain insights. This integrated approach strengthens the model's adaptability to individual user behaviors and results in more accurate, personalized recommendations. The final user representation is formulated as follows:

$$e^a_u = COM^a(e^a_{lu}, e^a_{agu}); e^b_u = COM^b(e^b_{lu}, e^b_{agu}), \tag{10}$$

where COM is a combination operation such as summation, averaging or an attention mechanism [29].

The final ranking score for a user on a target item is calculated by computing the dot product between the final user representation and the item representation. This score quantifies the relevance of the item to the user, enabling the model to prioritize items that best match the user's preferences.

4.5 Loss Function and Model Training

4.5.1 Loss Function. We compute the recommendation loss using the Bayesian personalized ranking loss [31]. For domain $\mathcal{D}^a$, the recommendation loss $\mathcal{L}^a_{rec}$ is calculated as follows:

$$\mathcal{L}^a_{\text{rec}} = \sum_{(u,m^a,n^a)\in O^a} -\log\sigma(\hat{r}^a_{um} - \hat{r}^a_{un}). \tag{11}$$

Similarly, for domain $\mathcal{D}^b$, the recommendation loss $\mathcal{L}^b_{rec}$ is

$$\mathcal{L}^b_{\text{rec}} = \sum_{(u,m^b,n^b)\in O^b} -\log\sigma(\hat{r}^b_{um} - \hat{r}^b_{un}), \tag{12}$$

where O^a represents the training data for domain $\mathcal{D}^a$, with $m^a \in O^{a+}$ denoting the set of interacted items and $n^a \in O^{a-}$ signifying the set of non-interacted items. Similarly, O^b refers to the training data in domain $\mathcal{D}^b$. The $\hat{r}$ terms denote the ranking score of a user to the target item.

The objective function for each domain is formulated as:

$$\begin{aligned} \mathcal{L}^a &= \mathcal{L}^a_{rec} + \lambda\mathcal{L}^a_{con} + \delta\|\Theta^a\|_2 \\ \mathcal{L}^b &= \mathcal{L}^b_{rec} + \lambda\mathcal{L}^b_{con} + \delta\|\Theta^b\|_2, \end{aligned} \tag{13}$$

where $\mathcal{L}^a_{con}$ and $\mathcal{L}^b_{con}$ represent the consistency losses calculated from Equation (2). The hyperparameter λ is used to control the influence of the consistency loss in each domain's objective function. Θ^* represents the set of parameters in each domain while the δ term regulates the strength of the $L2$-regularization, helping to prevent overfitting.

Algorithm 2: Training of FGD-CDR

```
Input  : $\mathcal{D}^a = (\mathcal{U}^a, \mathcal{I}^a, \mathcal{R}^a)$; $\mathcal{D}^b = (\mathcal{U}^b, \mathcal{I}^b, \mathcal{R}^b)$
         Hyperparameters $\lambda, \gamma, \sigma, k, \alpha, t_{com}$
Output : trained model
Random initialization;
while stopping criteria not met do
    Create negative samples and randomly shuffle the training data
    Domain $\mathcal{D}^a$ receives global-shared user representations $e^b_{gsu}$ from domain $\mathcal{D}^b$
    Domain $\mathcal{D}^b$ receives global-shared user representations $e^a_{gsu}$ from domain $\mathcal{D}^a$
    for $t = 1$ to $t_{com}$ do
        // In domain $\mathcal{D}^a$
        Extract the item, local user, global user representations:
        $e^a_i, e^a_{lu} \leftarrow LGE^a((e^a_i)^0, (e^a_{lu})^0, \mathsf{A}^a)$ ; $e^a_{gu} \leftarrow GGE^a((e^a_i)^0, (e^a_{gu})^0, \mathsf{A}^a)$
        Extract latent overlapping users: $\mathcal{I}C^{u^a}_{a \to b} \leftarrow CE^a(e^a_{gu}, e^b_{gsu}, k, \sigma)$;
        Generate adapted global user representations: $e^a_{agu} \leftarrow PPA^a(e^a_{gu}, \mathcal{I}C^{u^a}_{a \to b}, \alpha)$;
        Generate final user representations: $e^a_u \leftarrow COM^a(e^a_{lu}, e^a_{agu})$;
        Predict ranking scores: $\hat{r}^a_{ui} = {e^a_u}^T e^a_i$;
        Calculate objective function in domain $\mathcal{D}^a$ using Eq.2, 11, 12 and 13.
        Updates parameters related to domain $\mathcal{D}^a$

        // In domain $\mathcal{D}^b$
        Extract the item, local user, global user representations:
        $e^b_i, e^b_{lu} \leftarrow LGE^b((e^b_i)^0, (e^b_{lu})^0, \mathsf{A}^b)$ ; $e^b_{gu} \leftarrow GGE^b((e^b_i)^0, (e^b_{gu})^0, \mathsf{A}^b)$
        Extract latent overlapping users: $\mathcal{I}C^{u^b}_{b \to a} \leftarrow CE^b(e^b_{gu}, e^a_{gsu}, k, \sigma)$;
        Generate adapted global user representations: $e^b_{agu} \leftarrow PPA^b(e^b_{gu}, \mathcal{I}C^{u^b}_{b \to a}, \alpha)$;
        Generate final user representations: $e^b_u \leftarrow COM^b(e^b_{lu}, e^b_{agu})$;
        Predict ranking scores: $\hat{r}^b_{ui} = {e^b_u}^T e^b_i$;
        Calculate objective function in domain $\mathcal{D}^b$ using Eq.2, 11, 12 and 13.
        Updates parameters related to domain $\mathcal{D}^b$
    end
    Extract global-shared user representations in domain $\mathcal{D}^a$: $e^a_{gsu} \leftarrow GSGE^a((e^a_i)^0, (e^a_{gu})^0, \mathsf{A}^a, \epsilon)$
    Extract global-shared user representations in domain $\mathcal{D}^b$: $e^b_{gsu} \leftarrow GSGE^b((e^b_i)^0, (e^b_{gu})^0, \mathsf{A}^b, \epsilon)$
end
```

4.5.2 Model Training. As representations are exchanged across domains, communication costs can become a major bottleneck in federated learning environments, especially when handling large-scale data and frequent updates [30]. To mitigate this issue, we introduce an adjustable communication interval, denoted as t_{com}, which specifies the number of training epochs between cross-domain exchanges. Instead of sharing representations after every training iteration, domains perform multiple local updates and synchronize only after t_{com} epochs, thus reducing bandwidth usage and overall communication costs. In each domain, the process begins by receiving the latest global-shared user representations from other domains. The item, local user, and global user representations are then extracted using the LGE and the GGE. The model then leverages AKT to integrate information from multiple domains, make predictions, and optimize the objective function. After every t_{com} epoch, each domain extracts its updated global-shared user representations and shares them with other domains. The training process of our method in a two-domain setting is shown in Algorithm 2.

Regarding communication cost, suppose that a total of n_{round} local update rounds are executed, the number of global communication rounds is approximately n_{round}/t_{com}. Consequently, the overall communication cost is expressed as $O(d.n_{round}/t_{com}.n_u.|D|(|D|-1))$, where d is size of vector representations, n_u is the number of users, and $|D|$ is the number of domains. By performing global communication every t_{com} epochs, this strategy can achieve a t_{com}-fold reduction communication cost—the larger the t_{com}, the lower the communication overhead. However, setting t_{com} too high may result in the model converging to a suboptimal solution on local domains, thereby reducing the effectiveness of cross-domain knowledge transfer.

4.6 Algorithm Analysis

4.6.1 Analysis on Privacy Protection. In our federated learning framework, each business partner retains user data locally, aligning with Rule 1 that prohibits sharing raw interaction data between domains. To address Rule 2, which mandates the protection of user representations before exchange, we apply DP techniques, ensuring that transferred knowledge remains secure—not only from external attackers but also from other business partners participating in the knowledge transfer process. Furthermore, our transfer mechanism operates without relying on any user-identifying information, making it ideally suited for privacy-preserving settings where obtaining such identifying data is challenging.

Proof: The Algorithm 1 satisfies the edge-level DP.

Laplace achieves DP by adding noise drawn from Laplace distribution. The Laplace distribution is defined as:

$$Lap(\theta) = \frac{1}{2\theta} exp\left(\frac{-|x|}{\theta}\right), \tag{14}$$

where $\theta = \frac{\Delta_e}{\epsilon}$ is the noise scale, with Δ_e representing the sensitivity, x being a random variable that follows the Laplace distribution. For two neighboring graphs $\mathcal{G}$ and $\mathcal{G}'$ differing by one edge, the probability density function of the output O is:

$$\Pr[\mathcal{A}(\mathcal{G}) \in O] = \frac{1}{2\theta} exp\left(\frac{-|O - e_{\mathcal{G}}|}{\theta}\right), \tag{15}$$

and similarly for $\mathcal{G}'$

$$\Pr[\mathcal{A}(\mathcal{G}') \in O] = \frac{1}{2\theta} exp\left(\frac{-|O - e_{\mathcal{G}'}|}{\theta}\right), \tag{16}$$

where $e_{\mathcal{G}}$ and $e_{\mathcal{G}'}$ represents the node representations learned from graphs $\mathcal{G}$ and $\mathcal{G}'$, respectively.

Taking the ratio:

$$\frac{\Pr[\mathcal{A}(\mathcal{G}) \in O]}{\Pr[\mathcal{A}(\mathcal{G}') \in O]} = exp\left(\frac{|e_{\mathcal{G}} - e_{\mathcal{G}'}|}{\theta}\right). \tag{17}$$

Since $|e_{\mathcal{G}} - e_{\mathcal{G}'}| \leq \Delta_e$ and $\Delta_e = \theta\epsilon$, we obtain:

$$\frac{\Pr[\mathcal{A}(\mathcal{G}) \in O]}{\Pr[\mathcal{A}(\mathcal{G}') \in O]} \leq exp\left(\frac{\theta\epsilon}{\theta}\right) = exp(\epsilon). \tag{18}$$

4.6.2 Potential Extension Scenario Setting. Our transfer mechanism introduces a general mapping function between domains by extracting latent overlapping users based on general preferences. This enables our method to effectively adapt to scenarios where no shared users exist between domains. In real-world applications, this corresponds to collaborations among different companies, where user overlap may be non-existent or difficult to identify due to privacy concerns. Furthermore, our method is designed to support multi-domain scenarios, where each domain operates as an independent client. The training process of FGD-CDR in a multi-domain setting is thoroughly detailed in the Algorithm 3. To validate the effectiveness of our approach, we conducted experiments for both the non-overlapping user scenario and the multi-domain scenario. The results and analysis of these experiments are presented in Section 5.5.

Algorithm 3: Extension of FGD-CDR for Multi-domain Scenarios

```
Input  : $\mathcal{D} = \{\mathcal{D}^1, \mathcal{D}^2, \ldots, \mathcal{D}^n\}$, where each $\mathcal{D}^p = (\mathcal{U}^p, \mathcal{I}^p, \mathcal{R}^p)$ for $p = 1, 2, \ldots, n$
         $\mathcal{U}^p$: User set in domain $\mathcal{D}^p$
         $\mathcal{I}^p$: Item set in domain $\mathcal{D}^p$
         $\mathcal{R}^p$: Interaction matrix in domain $\mathcal{D}^p$
         Hyperparameters $\lambda$, $\gamma$, $\sigma$, $k$, $\alpha$, $t_{com}$
Output : trained model
Random initialization;
while stopping criteria not met do
    Create negative samples and randomly shuffle the training data
    for domain $\mathcal{D}^p \in \mathcal{D}$ do
        Received global-shared user representations from other domains
        for $t = 1$ to $t_{com}$ do
            Extract the item, local user, global user representations:
            $e_i^p, e_{lu}^p \leftarrow LGE^p((e_i^p)^0, (e_{lu}^p)^0, \mathsf{A}^p)$ ; $e_{gu}^p \leftarrow GGE^p((e_i^p)^0, (e_{gu}^p)^0, \mathsf{A}^p)$
            for pair of domains $(\mathcal{D}^p, \mathcal{D}^q)$, where $\mathcal{D}^p \neq \mathcal{D}^q$ do
                Extract latent overlapping users: $\mathcal{I}C_{p\to q}^{u^p} \leftarrow CE^p(e_{gu}^p, e_{gsu}^q, k, \sigma)$;
            end
            Generate adapted global user representations: $e_{agu}^p \leftarrow PPA^p(e_{gu}^p, \mathcal{I}C_{p\to cross}^{u^p}, \alpha)$;
                where $\mathcal{I}C_{p\to cross}^{u^p} = \{\mathcal{I}C_{p\to q}^{u^p} | \mathcal{D}^q \in \mathcal{D}, \mathcal{D}^p \neq \mathcal{D}^q\}$;
            Generate final user representations: $e_u^p \leftarrow COM^p(e_{lu}^p, e_{agu}^p)$;
            Predict ranking scores: $\hat{r}_{ui}^p = {e_u^p}^T e_i^p$;
            Calculate objective function in domain $\mathcal{D}^p$ using Eq.2, 11, 12 and 13.
            Updates parameters related to domain $\mathcal{D}^p$
        end
        Extract global-shared user representations: $e_{gsu}^p \leftarrow GSGE^p((e_i^p)^0, (e_{gu}^p)^0, \mathsf{A}^p, \epsilon)$
    end
end
```

5 Experiments

This section outlines the experiments conducted to assess the recommendation performance and privacy-preserving capabilities of our FGD-CDR model, targeting the following research questions:

- *RQ1*: How does FGD-CDR perform compared to state-of-the-art baseline models?
- *RQ2*: What is the impact of each individual component on the overall performance of FGD-CDR?
- *RQ3*: How effective is FGD-CDR in ensuring privacy protection?
- *RQ4*: Can FGD-CDR be adapted for other real-world scenarios?
- *RQ5*: How do hyperparameter choices affect the performance of FGD-CDR?

5.1 Experimental Settings

5.1.1 Datasets. We evaluate FGD-CDR, comparing it to other baseline models, using two widely adopted real-world datasets: Douban [35] and Amazon.[1] To ensure data quality, we filter out users with fewer than 10 interactions, focusing specifically on overlapping users across domains. For each user, two interactions per domain are randomly selected—one for validation and one for testing—while the remaining interactions form the training set. To simulate cross-domain recommendation tasks, we create three distinct tasks, with descriptive statistics for each provided in Table 2.

[1] https://jmcauley.ucsd.edu/data/amazon/.

Table 2. Statistics of Tasks

Task	Dataset	Users	Items	Ratings	Density
Task 1	amazon-movie (D1_1)	6,255	42,441	373,596	0.14%
	amazon-cd (D2_1)	6,255	57,695	309,692	0.09%
Task 2	amazon-cd (D2_2)	4,968	54,439	233,812	0.09%
	amazon-book (D3_1)	4,968	125,663	334,700	0.05%
Task 3	douban-music (D4_1)	7,106	16,813	571,738	0.48%
	douban-book (D5_1)	7,106	16,407	472,187	0.41%

5.1.2 Evaluation Protocol. For each record in the evaluation data, we randomly select 99 items that the user has not interacted with to generate negative samples. The model then predicts preferences for 100 items, consisting of 1 positive sample and 99 negative samples, and produces a ranked list of the top-N items. To assess performance, we use the standard metrics: hit ratio ($HR@N$) and normalized discounted cumulative gain ($NDCG@N$), setting N to 5. Each experiment is repeated five times to ensure robust results.

5.1.3 Baselines. To compare the performance of our proposed approach, we have selected the following baseline methods:

Single-domain Recommendation Methods:

- *LightGCN* [11]: captures high-order connectivity by propagating and aggregating information within the user-item interaction graph.
- *FCF* [1]: employs federate learning that trains local models on each client and aggregates local parameters to create a global model.
- *impactOfDP* [28]: applies DP to protect interaction data at the graph level.

Non-private Cross-domain Recommendation Methods:

- *PPGN* [44]: builds a cross-domain graph to establish connections among domains and exploits this graph to learn representations.
- *Bi-TGCF* [21]: performs the bidirectional knowledge transfer at each graph convolution layer.
- *DisenCDR* [2]: utilizes disentangled learning to separate domain-shared and domain-specific information.

Privacy-preserving Cross-domain Recommendation Methods:

- *PriCDR* [4]: implements a DP algorithm to protect the source rating matrix before transfer to the target domain.
- *P2FCDR* [5]: develops an orthogonal mapping matrix to learn and protect global user representations with DP before transferring across domains.
- *PPCDR* [37]: designs a graph transfer module that fuses local and global user representations during private updates and applies DP to protect global user representations before transfer.

5.1.4 Implementation and Hyperparameter Settings. Common hyperparameters are set across all methods as follows: the embedding size is 64, the batch size is 1,024, the learning rate is 0.001, the L2 regularization coefficient is 1e-5, three graph convolution layers, and the Adam optimizer to learn the parameters. The maximum number of training epochs is set to 200. Specific parameters for each baseline are set based on the original papers and fine-tuned accordingly. Our method is implemented in PyTorch, with the key hyperparameters λ, γ, σ, α, k, and t_{com} tune via a grid search. We apply the same privacy budget across all graph layers.

Table 3. Overall Performance Comparison of FGD-CDR with Baseline Methods for All Tasks

Method	Task 1				Task 2				Task 3			
	D1_1		D2_1		D2_2		D3_1		D4_1		D5_1	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
LightGCN	0.5243	0.3745	0.5189	0.3841	0.4401	0.3172	0.3615	0.2644	0.6191	0.4669	0.5876	0.4403
FCF	0.4109	0.2859	0.4091	0.2987	0.3492	0.2527	0.2874	0.2109	0.5741	0.4265	0.5276	0.3916
impactOfDP	0.4814	0.3464	0.4719	0.3538	0.4147	0.3087	0.3422	0.2577	0.5913	0.4455	0.5586	0.4175
PPGN	0.5550	0.4098	0.5239	0.3970	0.4617	0.3416	0.3957	0.3058	0.6492	0.4930	0.6442	0.4913
Bi-TGCF	0.5646	0.4151	0.5421	0.4016	0.4801	0.3541	0.4101	0.3068	0.6711	0.5178	0.6538	0.5064
DisenCDR	0.5741	0.4204	0.5656	0.4266	0.4915	0.3661	0.4092	0.3127	0.6701	0.5108	0.6469	0.4919
PriCDR	0.5183	0.3765	0.5249	0.3922	0.4613	0.3426	0.3861	0.2914	0.6365	0.4803	0.6033	0.4512
P2FCDR	0.5380	0.3878	0.4903	0.3551	0.4322	0.3157	0.3570	0.2705	0.6546	0.4964	0.6245	0.4717
PPCDR	0.5548	0.4084	0.5645	0.3896	0.4905	0.3564	0.4137	0.3097	0.6718	0.5185	0.6552	0.5064
FGD-CDR	**0.6312**	**0.4853**	**0.6293**	**0.4943**	**0.5539**	**0.4321**	**0.4571**	**0.3646**	**0.7306**	**0.5827**	**0.7073**	**0.5628**
RI	9.9%	15.4%	11.3%	15.9%	12.7%	18.0%	10.5%	16.6%	8.7%	12.4%	8.0%	11.1%

5.2 Overall Performance Comparison (RQ1)

Table 3 presents the results for *HR*@5 and *NDCG*@5, comparing our method against selected baseline methods across the three tasks. The privacy budget used for our FGD-CDR and privacy-preserving baseline methods is set to 1.0. The highest performances are highlighted in bold, while the top-performing baseline results are underlined. "RI" represents the relative improvement of our FGD-CDR model compared to the best-performing baseline.

The insights gained from these experimental results are as follows:

- GNNs are increasingly preferred in modern recommendation approaches due to their ability to capture intricate collaborative filtering signals from user-item interaction graphs. Utilizing advanced propagation and aggregation mechanisms, GNNs excel at modeling complex relationships. The superior performance of graph-based methods, such as PPCDR, over traditional neural network-based methods like PriCDR and P2FCDR highlights this advantage.
- Cross-domain recommendation methods typically outperform single-domain approaches, demonstrating the benefits of sharing valuable information across domains to mitigate data sparsity challenges.
- PPGN relies solely on global user representations, limiting its ability to capture nuanced user preferences. In contrast, methods that integrate both local and global user representations—such as Bi-TGCF, DisenCDR, and our FGD-CDR—achieve superior performance. However, Bi-TGCF lacks a mechanism to effectively distinguish between local and global preferences, which can lead to negative transfer when domain-specific information is embedded in global user representations. Additionally, as highlighted in [45], physically overlapping users may exhibit different preferences across domains. Consequently, bidirectional knowledge transfer methods based on physical overlap, such as Bi-TGCF and DisenCDR, may introduce noise, particularly when these users display conflicting preferences in different domains. In contrast, our method introduces consistency loss, which effectively differentiates local and global user representations, preventing them from collapsing into an indistinguishable space. Furthermore, instead of relying on physical overlap, our approach identifies latent overlapping users based on similarities in general preferences. This ensures that users with aligned preferences are meaningfully connected, reducing noise and enhancing the effectiveness of knowledge transfer across domains.

Table 4. Results of Ablation Studies in All Tasks

Method	Task 1				Task 2				Task 3			
	D1_1		D2_1		D2_2		D3_1		D4_1		D5_1	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
w/o sim	0.5681	0.4369	0.5552	0.4452	0.4830	0.3823	0.4000	0.3216	0.7030	0.5490	0.6803	0.5296
w/o ppa	0.5796	0.4380	0.6000	0.4671	0.5212	0.4013	0.4340	0.3422	0.7081	0.5536	0.6878	0.5422
w/o con-loss	0.5816	0.4321	0.5822	0.4447	0.5001	0.3759	0.4184	0.3228	0.6796	0.5202	0.6451	0.4931
FGD-CDR	**0.6312**	**0.4853**	**0.6293**	**0.4943**	**0.5539**	**0.4321**	**0.4571**	**0.3646**	**0.7306**	**0.5827**	**0.7073**	**0.5628**

The highest performances are highlighted in bold.

- The impactOfDP method applies DP at the graph-data level to safeguard the structure of the interaction graph. However, the notable performance drop compared to LightGCN in Task 1 suggests that applying DP may distort key structural elements of the graph, thereby impacting recommendation accuracy.
- PriCDR improves recommendation performance in sparser domains through unidirectional transfer. However, it may be less effective when transferring data from sparser to richer domains. For example, in Task 1, PriCDR better enhances the recommendations in D2_1 (sparser) compared to the single-domain method LightGCN but performs worse in D1_1. This observation underscores the limitations of unidirectional transfer methods. P2FCDR employs bidirectional transfer, showing better performance than PriCDR in Tasks 1 and 3. However, it utilizes deep matrix factorization as the encoder, which constrains its ability to capture high-quality representations.
- Among privacy-preserving baseline methods, PPCDR demonstrates better performance by effectively extracting both local and global user preferences. However, similar to Bi-TGCF, it lacks a mechanism to distinguish between these preferences, potentially leading to negative transfer. Furthermore, those privacy-preserving CDRS require user-identifying information to detect overlapping users for cross-domain connections, limiting their practical deployment, especially in high-security scenarios.
- Our proposed method, FGD-CDR, consistently outperforms all baseline methods, highlighting its effectiveness in handling data sparsity, ensuring privacy protection, and preventing negative transfer. Specifically, FGD-CDR improves performance by an average of 35% compared to single-domain recommendation methods and by 14% when compared to non-private CDRS. For privacy-preserving CDRS, FGD-CDR shows an improvement of around 20% compared to PriCDR and P2FCDR, and approximately 12% compared to PPCDR.

5.3 Model Ablation Study (RQ2)

We perform ablation studies to evaluate the impact of key components in FGD-CDR. These experiments include: (1) a similarity computation step to select the top-k most similar users for establishing cross-domain connections (*w/o sim*), (2) PPA for generating adapted global user representations (*w/o ppa*), and (3) the application of a consistency loss (*w/o con-loss*). As shown in Table 4, removing any of these components leads to a decline in performance, highlighting their critical role in our model.

First, the *w/o sim* variant replaces the similarity computation with random selection for the top-k user selection. This modification leads to a 13% performance drop, illustrating the importance of similarity-based connections. By matching users with similar behaviors across domains, the model efficiently transfers knowledge, improving recommendation performance. However, randomly pairing users with unrelated behaviors introduces noise into the knowledge transfer process, significantly weakening the model's ability to generate accurate recommendations.

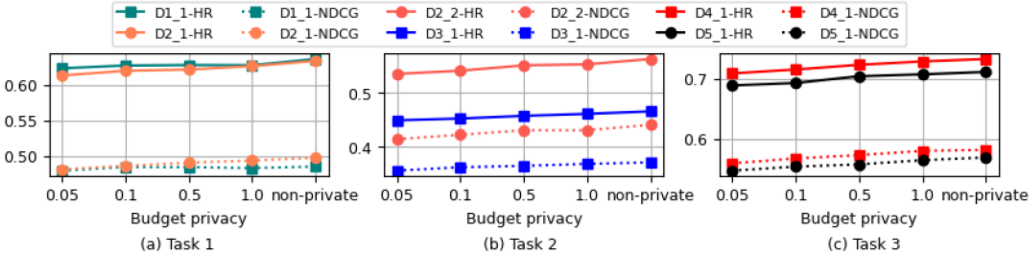

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 2. Impact of privacy budget ϵ on recommendation performance in our FGD-CDR.

Second, in the *w/o ppa* variant, the aggregation function in Equation (8) is replaced with simple averaging, where the adapted global user representation is calculated as the average of the global user representation and the top-k global-shared representations from other domains. This change results in a 6% performance decline, highlighting the value of personalized aggregation. Unlike simple averaging, the personalized aggregation mechanism controls the integration of cross-domain knowledge, filtering out irrelevant information while preserving critical knowledge during the transfer. This ensures that the adapted global user representation is not diluted by noisy or unrelated domain data.

Third, removing the consistency loss in the *w/o con-loss* variant results in a 10% performance decline, emphasizing its role in representation learning. Consistency loss enhances alignment between local and global representations of the same user, preventing them from being collapsed. This mechanism creates a balanced representation space, ensuring that both shared and domain-specific information are well-represented, thus reducing the risk of irrelevant knowledge transfer between domains.

In summary, the results of these ablation studies confirm that each of these components—similarity computation, PPA, and consistency loss—plays a pivotal role in the effectiveness of FGD-CDR. Their removal disrupts the balance of knowledge transfer, degrades representation learning, and ultimately reduces the model's overall effectiveness in cross-domain recommendation tasks.

5.4 Privacy Analysis (RQ3)

We also evaluate the impact of privacy mechanisms in our FGD-CDR model on both recommendation performance and privacy protection. First, we compare our method across different privacy budget settings with a non-private variant, representing a scenario in which no privacy mechanism is applied to the aggregation outputs. Following the experiments in [37], we assess the privacy protection capability of FGD-CDR by predicting historical interactions based on global-shared user representations. For each interaction record, we randomly sample nine non-interacted items and combine them with the interacted item to form a candidate set. We then rank all candidate items based on a score computed as the dot product between the item representation and the global-shared user representations. The item with the highest score is predicted as the historical interaction. Lower accuracy in this task indicates better privacy protection.

The results for overall recommendation performance and privacy protection capabilities are presented in Figures 2 and 3, respectively. As shown, an increase in the privacy budget, denoted by ϵ, leads to improved recommendation performance but decreased privacy protection, and vice versa. A smaller ϵ necessitates the addition of more noise to the representations, which enhances privacy guarantees but hampers the model's ability to accurately capture user preferences, resulting in reduced recommendation accuracy. In contrast, a larger ϵ improves recommendation accuracy but compromises privacy protection. This highlights the inherent tradeoff between privacy and accuracy.

Next, we compare our method with selected cross-domain baselines, including a non-private method (Bi-TGCF) and two privacy-preserving methods (P2FCDR and PPCDR). The results are

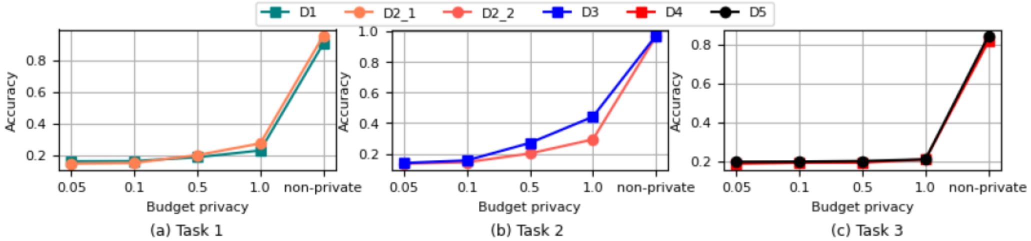

Fig. 3. Impact of privacy budget ϵ on privacy protection ability in our FGD-CDR. The lower accuracy reflects stronger privacy protection.

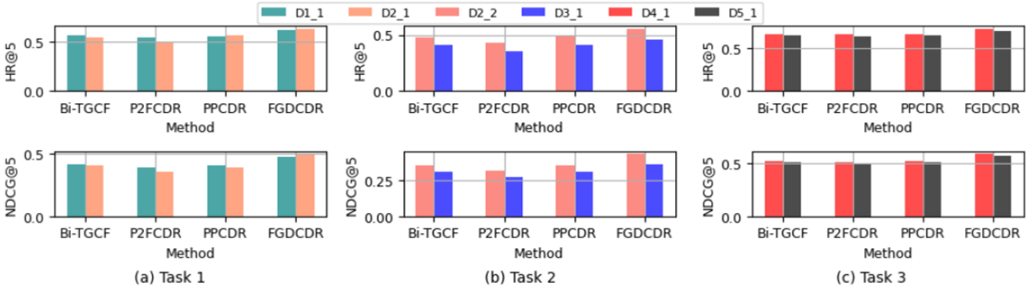

Fig. 4. Comparison of recommendation performance between our FGD-CDR and selected baseline methods.

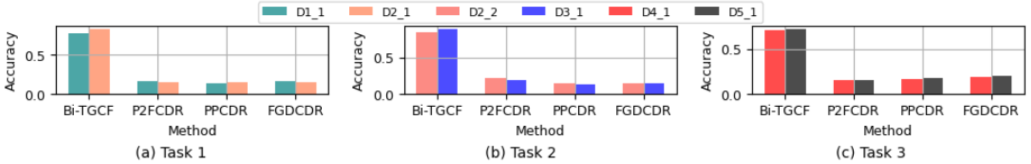

Fig. 5. Comparison in privacy protection ability between our FGD-CDR and selected baseline methods. The lower accuracy reflects stronger privacy protection.

shown in Figures 4 and 5. The non-private method, Bi-TGCF, demonstrates performance comparable to the privacy-preserving method PPCDR; however, it carries the highest risk of inferring user historical data. Our FGD-CDR model achieves superior recommendation accuracy, while PPCDR offers stronger privacy protection by applying DP to global user representations. This method adds noise at the output layer, which can overwhelm the learned representations, enhancing privacy protection but often compromising the model's ability to retain useful information and thus reducing recommendation accuracy. In contrast, our method introduces noise at each aggregation step during graph representation learning. By incrementally adding noise after each aggregation, our model preserves more informative signals at each layer, enabling it to refine representations in subsequent layers while smoothing out the noise from earlier steps. This approach results in better accuracy, albeit with slightly weaker privacy protection. As demonstrated in Figures 4 and 5, our method achieves a better balance between maintaining high recommendation accuracy and delivering effective privacy protection.

5.5 Further Analysis (RQ4)

This section assesses the effectiveness of our approach in non-overlapping and multi-domain scenarios, with the tasks are defined in Table 5.

Table 5. Statistics of Tasks in Non-overlapping and Multi-domain Scenarios

Task	Category	Users	Items	Ratings	Density
Non-overlapping Cross-domain Recommendation Tasks					
Task 4	douban-music (D4_2)	10,393	11,589	1,143,883	0.95%
	amazon-cd (D2_3)	7,062	27,456	378,391	0.20%
Task 5	douban-book (D5_2)	8,821	8,627	808,298	1.06%
	amazon-book (D3_2)	9,407	20,058	643,397	0.34%
Multi-domain Recommendation Tasks					
Task 6	amazon-movie (D1_2)	4,327	44,798	313,815	0.16%
	amazon-cd (D2_4)	4,327	60,290	216,259	0.08%
	amazon-book (D3_3)	4,327	139,268	343,488	0.06%
Task 7	douban-music (D4_2)	10,393	11,589	1,143,883	0.95%
	douban-book (D5_2)	8,821	8,627	808,298	1.06%
	amazon-book (D3_2)	9,407	20,058	643,397	0.34%

Table 6. Overall Performance Comparison for Non-overlapping Tasks

Method	Task 4				Task 5			
	D4_2		D2_3		D5_2		D3_2	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
LightGCN	0.5938	0.4406	0.5930	0.4340	0.5622	0.4178	0.5290	0.3777
FCF	0.5425	0.3972	0.5290	0.3899	0.5119	0.3751	0.4555	0.3269
impactOfDP	0.5731	0.4252	0.5571	0.4105	0.5477	0.4061	0.4811	0.3436
FGD-CDR	**0.6725**	**0.5166**	**0.7183**	**0.5624**	**0.6510**	**0.5010**	**0.6498**	**0.4845**

The highest performances are highlighted in bold, while the top-performing baseline results are underlined.

Table 7. Overall Performance Comparison for Multi-domain Tasks

Method	Task 6						Task 7					
	D1_2		D2_4		D3_3		D4_2		D5_2		D3_2	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
LightGCN	0.4635	0.3403	0.4007	0.2901	0.3463	0.2522	0.5938	0.4406	0.5622	0.4178	0.5290	0.3777
FCF	0.3460	0.2345	0.2911	0.2091	0.2569	0.1901	0.5425	0.3972	0.5119	0.3751	0.4555	0.3269
impactOfDP	0.4234	0.3004	0.3790	0.2813	0.3082	0.2267	0.5731	0.4252	0.5477	0.4061	0.4811	0.3436
PPGN	0.5038	0.3663	0.4294	0.3249	0.3626	0.2704	-	-	-	-	-	-
PPCDR	0.5137	0.3726	0.4638	0.3356	0.3854	0.2843	-	-	-	-	-	-
FGPCDR	**0.5735**	**0.4339**	**0.5098**	**0.4008**	**0.4313**	**0.3397**	**0.6829**	**0.5261**	**0.6583**	**0.5053**	**0.6601**	**0.4936**

The highest performances are highlighted in bold, while the top-performing baseline results are underlined.

Table 6 provides a detailed comparison of recommendation performance between our FGD-CDR model and baseline methods in non-overlapping scenarios. Our model demonstrates a substantial improvement over single-domain approaches, underscoring the effectiveness of our transfer mechanism. This enhancement can be attributed to two key factors. First, the implicit relationships created through latent overlapping user help to limit information transfer from users who are exclusively interested in items within the current domain. Second, more precise similarity calculations lower the chances of creating inaccurate implicit connections.

Table 7 illustrates the overall recommendation performance for the multi-domain tasks. Notably, the performance in each domain for Task 7 is superior to its corresponding performance in Task 4 (D4_2) and Task 5 (D5_2 and D3_2), with an average improvement of

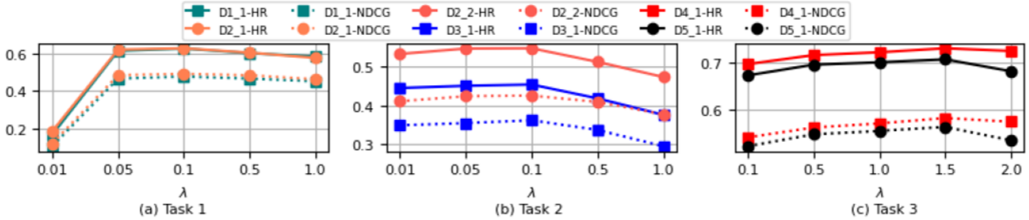

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 6. Hyperparameter analysis of λ.

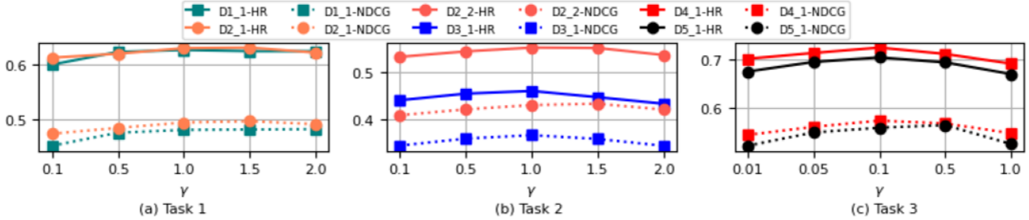

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 7. Hyperparameter analysis of γ.

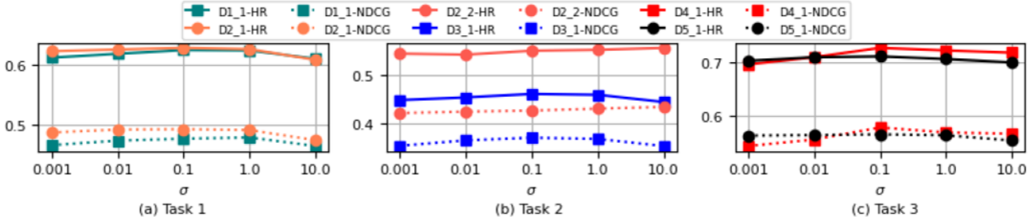

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 8. Hyperparameter analysis of σ.

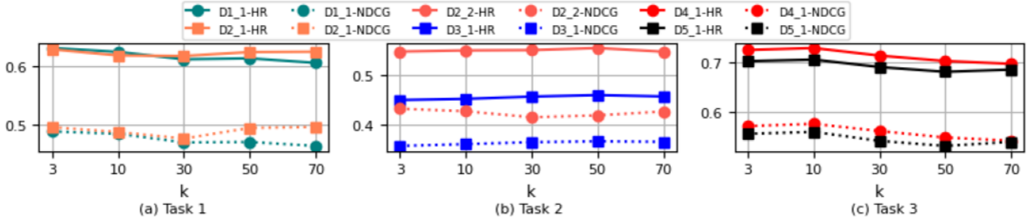

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 9. Hyperparameter analysis of k.

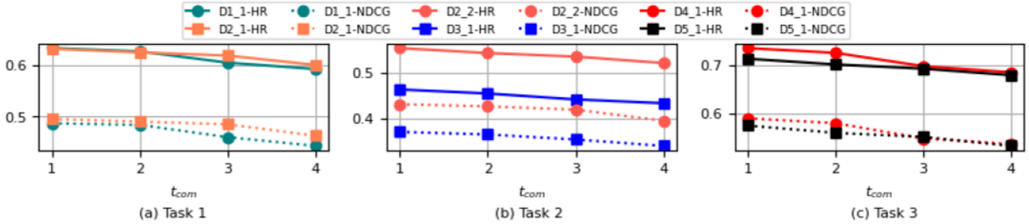

(a) Task 1 (b) Task 2 (c) Task 3

Fig. 10. Hyperparameter analysis of t_{com}.

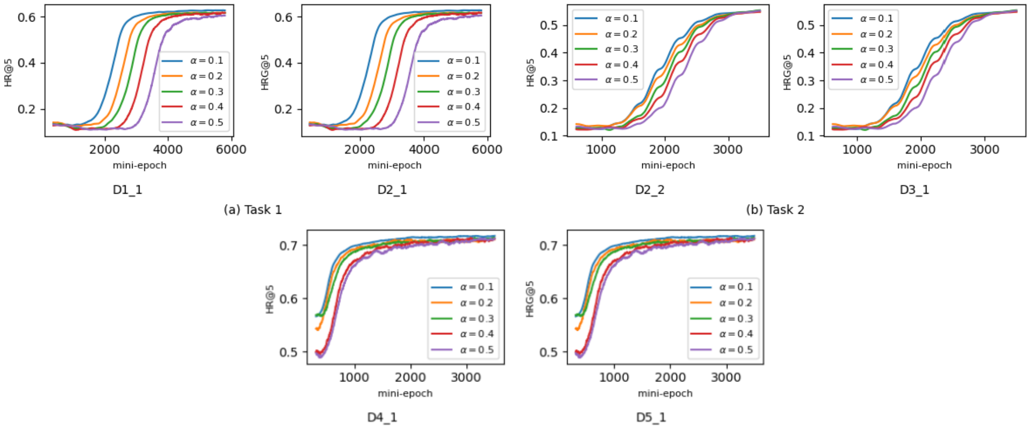

D1_1 D2_1

(a) Task 1

D2_2 D3_1

(b) Task 2

D4_1 D5_1

(c) Task 3

Fig. 11. Hyperparameter analysis of α.

around 2%. This improvement highlights the advantage of leveraging knowledge from multiple domains to enhance recommendations within each individual domain.

5.6 Hyperparameter Analysis (RQ5)

This section analyzes the sensitivity of key hyperparameters λ, γ, σ, k, t_{com}, and α on the model's performance, as illustrated in the Figures 6–11. The results show that λ, γ, and t_{com} have a significant impact on the model's performance, while k has a lesser effect. Both α and σ exhibit stable performance across their respective ranges.

- λ controls the weight of the consistency loss relative to the recommendation loss. As the consistency loss promotes coherence between local and global user representations, increasing λ may help enforce robustness across domains, allowing the model to leverage cross-domain information. However, placing too much emphasis on this consistency can lead to overly homogeneous representations, potentially compromising the model's ability to optimize for specific domain preferences in recommendations. Conversely, reducing λ allows the model to focus more heavily on the recommendation loss. A small λ may lead to a model that performs well in individual domains but could suffer in cross-domain generalization due to less consistency across representations. The model achieves the highest performance at $\lambda = 0.1$ for Tasks 1 and 2, while for Task 3, the optimal value is $\lambda = 1.5$.
- The uniformity term, controlled by γ, ensures that local and global user representations remain diverse and well-distributed, preventing them from becoming overly similar. A higher γ encourages more distinct and varied representations, helping to reduce the risk of domain-specific information dominating the global representations, thus minimizing the likelihood of negative transfer. On the other hand, a lower γ enforces tighter alignment between local and global representations, which can enhance cross-domain generalization but might reduce the model's ability to capture nuanced, domain-specific user preferences. The model achieves optimal performance in Tasks 1 and 2 with $\gamma = 1.0$, while in Task 3, the best results occur with $\gamma = 0.1$.
- σ exhibits consistently strong performance in terms of $HR@5$ and $NDCG@5$, with stable results as σ varies between 0.001 and 10. This indicates that σ has less sensitivity compared to λ and γ, contributing to stable recommendation performance across different values.

- k determines the number of most similar users used to establish connections between domains. Increasing k results in denser implicit connections, facilitating greater cross-domain information exchange. However, this can also introduce noise if inaccurate connections are formed. Moreover, bidirectional knowledge transfer may not be symmetric, meaning the optimal value of k could vary across domains. The model performs optimally with $k = 3$ in Task 1, $k = 50$ in Task 2, and $k = 10$ in Task 3.
- t_{com} denotes the number of training epochs a domain performs before sharing its user preferences with other domains. Increasing t_{com} can help lower communication costs by reducing the frequency of data exchange. However, setting t_{com} too high risks degrading recommendation performance, as the model might converge to suboptimal solutions based on local domain data. To maintain an effective balance between communication efficiency and recommendation accuracy, it is crucial to carefully select an appropriate t_{com} value. In our experiments, we found that setting t_{com} to 2 achieves a good compromise across all tasks, minimizing communication overhead while still providing strong recommendation performance.
- α defines the transfer ratio, indicating how much knowledge from other domains is incorporated into the current domain. While the convergence curves exhibit slight variations, they ultimately reach similar performance levels, demonstrating the model's robustness. As α increases, the learning curve initially grows more slowly, as the influence of knowledge from other domains may temporarily hinder the learning process in the current domain. However, after a few iterations, this external knowledge begins to support the target domain, leading to improved performance in subsequent stages. Generally, $\alpha = 0.3$ is an effective choice for the transfer ratio.

6 Conclusion

In this work, we propose a novel privacy-preserving CDRS that leverages federated transfer learning and DP. Our development process begin by designing a consistency mechanism to ensure the local and global user representations capture coherent information while, at the same time, preventing their collapse. Next, we implement privacy mechanisms, applying DP to the outputs of each aggregation step in the GNN. This effectively balances privacy and accuracy. We then introduce a general and stable transfer mechanism that extracts latent overlapping users to establish implicit connections, incorporating transferred knowledge through PPA. Experiments conducted on real-world datasets highlight the effectiveness of our proposed FGD-CDR, which outperforms baseline methods in terms of accuracy while also achieving superior privacy protection.

Future work will focus on integrating user profiles to gain deeper insights into user preferences. However, incorporating sensitive user information poses challenges in maintaining strong privacy guarantees while ensuring accurate personalized recommendations. Additionally, since deeper GNNs may experience significant privacy budget compounding, we will explore adaptive privacy budgeting as a potential solution.

References

[1] Muhammad Ammad-Ud-Din, Elena Ivannikova, Suleiman A. Khan, Were Oyomno, Qiang Fu, Kuan Eeik Tan, and Adrian Flanagan. 2019. Federated collaborative filtering for privacy-preserving personalized recommendation system. arXiv:1901.09888. Retrieved from https://arxiv.org/abs/1901.09888

[2] Jiangxia Cao, Xixun Lin, Xin Cong, Jing Ya, Tingwen Liu, and Bin Wang. 2022. Disencdr: Learning disentangled representations for cross-domain recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 267–277.

[3] Di Chai, Leye Wang, Kai Chen, and Qiang Yang. 2020. Secure federated matrix factorization. *IEEE Intelligent Systems* 36, 5 (2020), 11–20.

[4] Chaochao Chen, Huiwen Wu, Jiajie Su, Lingjuan Lyu, Xiaolin Zheng, and Li Wang. 2022. Differential private knowledge transfer for privacy-preserving cross-domain recommendation. In *Proceedings of the ACM Web Conference 2022*, 1455–1465.

[5] Gaode Chen, Xinghua Zhang, Yijun Su, Yantong Lai, Ji Xiang, Junbo Zhang, and Yu Zheng. 2023. Win-win: A privacy-preserving federated framework for dual-target cross-domain recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37, 4149–4156.

[6] Qiang Cui, Tao Wei, Yafeng Zhang, and Qing Zhang. 2020. HeroGRAPH: A heterogeneous graph framework for multi-target cross-domain recommendation. In *ORSUM@ RecSys*.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference (TCC '06)*. Springer, 265–284.

[8] Hamid Ebadi, David Sands, and Gerardo Schneider. 2015. Differential privacy: Now it's getting personal. *ACM SIGPLAN Notices* 50, 1 (2015), 69–81.

[9] Chen Gao, Xiangning Chen, Fuli Feng, Kai Zhao, Xiangnan He, Yong Li, and Depeng Jin. 2019. Cross-domain recommendation without sharing user-relevant data. In *The World Wide Web Conference*, 491–502.

[10] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. 2009. Accurate estimation of the degree distribution of private networks. In *Proceedings of the 2009 9th IEEE International Conference on Data Mining*. IEEE, 169–178.

[11] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 639–648.

[12] Xi He and Shufan Zhang. 2023. Differential privacy with fine-grained provenance: Opportunities and challenges. *Data Engineering* 47, 2 (2023), 21.

[13] Guangneng Hu, Yu Zhang, and Qiang Yang. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 667–676.

[14] Wenke Huang, Mang Ye, Zekun Shi, Guancheng Wan, He Li, Bo Du, and Qiang Yang. 2024. Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 12 (2024), 9387–9406.

[15] Danish Javeed, Muhammad Shahid Saeed, Prabhat Kumar, Alireza Jolfaei, Shareeful Islam, and A. K. M. Najmul Islam. 2023. Federated learning-based personalized recommendation systems: An overview on security and privacy challenges. *IEEE Transactions on Consumer Electronics* 70, 1 (2023), 2618–2627.

[16] Meng Jian, Yulong Bai, Xusong Fu, Jingjing Guo, Ge Shi, and Lifang Wu. 2024. Counterfactual graph convolutional learning for personalized recommendation. *ACM Transactions on Intelligent Systems and Technology* 15, 4 (2024), 1–20.

[17] Pan Li, Brian Brost, and Alexander Tuzhilin. 2022. Adversarial learning for cross domain recommendations. *ACM Transactions on Intelligent Systems and Technology* 14, 1 (2022), 1–25.

[18] Xinting Liao, Weiming Liu, Xiaolin Zheng, Binhui Yao, and Chaochao Chen. 2023. Ppgencdr: A stable and robust framework for privacy-preserving cross-domain recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37, 4453–4461.

[19] Dianqi Liu, Liang Bai, Tianyuan Yu, and Aiming Zhang. 2022. Towards method of horizontal federated learning: A survey. In *Proceedings of the 2022 8th International Conference on Big Data and Information Analytics (bigdia)*. IEEE, 259–266.

[20] Huiting Liu, Lingling Guo, Peipei Li, Peng Zhao, and Xindong Wu. 2021. Collaborative filtering with a deep adversarial and attention network for cross-domain recommendation. *Information Sciences* 565 (2021), 370–389.

[21] Meng Liu, Jianjun Li, Guohui Li, and Peng Pan. 2020. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 885–894.

[22] Yang Liu, Yan Kang, Tianyuan Zou, Yanhong Pu, Yuanqin He, Xiaozhou Ye, Ye Ouyang, Ya-Qin Zhang, and Qiang Yang. 2024. Vertical federated learning: Concepts, advances, and challenges. *IEEE Transactions on Knowledge and Data Engineering* 36, 7 (2024), 3615–3634.

[23] Tong Man, Huawei Shen, Xiaolong Jin, and Xueqi Cheng. 2017. Cross-domain recommendation: An embedding and mapping approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, Vol. 17, 2464–2470.

[24] Mingsong Mao, Jie Lu, Jialin Han, and Guangquan Zhang. 2019. Multiobjective e-commerce recommendations based on hypergraph ranking. *Information Sciences* 471 (2019), 269–287.

[25] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*. PMLR, 1273–1282.

[26] Wu Meihan, Li Li, Chang Tao, Eric Rigall, Wang Xiaodong, and Xu Cheng-Zhong. 2022. Fedcdr: Federated cross-domain recommendation for privacy-preserving rating prediction. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2179–2188.

[27] Peter Müllner, Elisabeth Lex, Markus Schedl, and Dominik Kowald. 2023. ReuseKNN: Neighborhood reuse for differentially private KNN-based recommendations. *ACM Transactions on Intelligent Systems and Technology* 14, 5 (2023), 1–29.

[28] Peter Müllner, Elisabeth Lex, Markus Schedl, and Dominik Kowald. 2024. The impact of differential privacy on recommendation accuracy and popularity bias. In *Proceedings of the European Conference on Information Retrieval*. Springer, 466–482.

[29] Zhaoyang Niu, Guoqiang Zhong, and Hui Yu. 2021. A review on the attention mechanism of deep learning. *Neurocomputing* 452 (2021), 48–62.

[30] Yang Pei, Renxin Mao, Yang Liu, Chaoran Chen, Shifeng Xu, Feng Qiang, and Blue Elephant Tech. 2021. Decentralized federated graph neural networks. In *Proceedings of the International Workshop on Federated and Transfer Learning for Data Sparsity and Confidentiality in Conjunction with IJCAI*.

[31] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 452–461.

[32] Sudipan Saha and Tahir Ahmad. 2021. Federated transfer learning: Concept and applications. *Intelligenza Artificiale* 15, 1 (2021), 35–44.

[33] Sina Sajadmanesh, Ali Shahin Shamsabadi, Aurélien Bellet, and Daniel Gatica-Perez. 2023. {GAP}: Differentially private graph neural networks with aggregation perturbation. In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)*, 3223–3240.

[34] Qusai Shambour and Jie Lu. 2015. An effective recommender system by unifying user and item trust information for B2B applications. *Journal of Computer and System Sciences* 81, 7 (2015), 1110–1126.

[35] Weiping Song, Zhiping Xiao, Yifan Wang, Laurent Charlin, Ming Zhang, and Jian Tang. 2019. Session-based social recommendation via dynamic graph attention networks. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining*, 555–563.

[36] Zehua Sun, Yonghui Xu, Yong Liu, Wei He, Lanju Kong, Fangzhao Wu, Yali Jiang, and Lizhen Cui. 2024. A survey on federated recommendation systems. *IEEE Transactions on Neural Networks and Learning Systems* 36, 1 (2024), 6–20.

[37] Changxin Tian, Yuexiang Xie, Xu Chen, Yaliang Li, and Xin Zhao. 2024. Privacy-preserving cross-domain recommendation with federated graph learning. *ACM Transactions on Information Systems* 42, 5 (2024), 1–29.

[38] Antonia Vlahou, Dara Hallinan, Rolf Apweiler, Angel Argiles, Joachim Beige, Ariela Benigni, Rainer Bischoff, Peter C. Black, Franziska Boehm, Jocelyn Céraline, et al. 2021. Data sharing under the general data protection regulation: Time to harmonize law and research ethics? *Hypertension* 77, 4 (2021), 1029–1035.

[39] Chuhan Wu, Fangzhao Wu, Lingjuan Lyu, Yongfeng Huang, and Xing Xie. 2022. FedCTR: Federated native ad CTR prediction with cross-platform user behavior data. *ACM Transactions on Intelligent Systems and Technology (TIST)* 13, 4 (2022), 1–19.

[40] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S. Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2020), 4–24.

[41] Mengmeng Yang, Taolin Guo, Tianqing Zhu, Ivan Tjuawinata, Jun Zhao, and Kwok-Yan Lam. 2023. Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces* 89 (2023), 103827.

[42] Ruiping Yin, Kan Li, Guangquan Zhang, and Jie Lu. 2019. A deeper graph neural network for recommender systems. *Knowledge-Based Systems* 185 (2019), 105020.

[43] Qian Zhang, Wenhui Liao, Guangquan Zhang, Bo Yuan, and Jie Lu. 2021. A deep dual adversarial network for cross-domain recommendation. *IEEE Transactions on Knowledge and Data Engineering* 35, 4 (2021), 3266–3278.

[44] Cheng Zhao, Chenliang Li, and Cong Fu. 2019. Cross-domain recommendation via preference propagation graphnet. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2165–2168.

[45] Yi Zhao, Chaozhuo Li, Jiquan Peng, Xiaohan Fang, Feiran Huang, Senzhang Wang, Xing Xie, and Jibing Gong. 2023. Beyond the overlapping users: cross-domain recommendation via adaptive anchor link learning. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1488–1497.

[46] Feng Zhu, Yan Wang, Chaochao Chen, Guanfeng Liu, and Xiaolin Zheng. 2020. A graphical and attentional framework for dual-target cross-domain recommendation. In *Proceedings of the 29th International Conference on International Joint Conferences on Artificial Intelligence*, 3001–3008.

[47] Feng Zhu, Yan Wang, Chaochao Chen, Jun Zhou, Longfei Li, and Guanfeng Liu. 2021. Cross-domain recommendation: Challenges, progress, and prospects. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI '21)*. International Joint Conferences on Artificial Intelligence, 4721–4728.

Received 13 November 2024; revised 7 April 2025; accepted 5 May 2025