

Article

Privacy-Preserving Peer-to-Peer Cross-Domain Collaborative Filtering via Intent-Adaptive Graph Reconstruction

Munan Li ¹ , Hao Zhang ¹, Jialong Li ² and Sinan Chen ^{3,4,*} 

¹ School of Artificial Intelligence, Dalian Maritime University, Dalian 116026, China; limunan@dlmu.edu.cn (M.L.); zh1120252238@dlmu.edu.cn (H.Z.)

² Waseda Institute for Advanced Study, Waseda University, Tokyo 162-0052, Japan; lijialong@fuji.waseda.jp

³ Center of Mathematical and Data Sciences, Kobe University, Kobe 657-8501, Japan

⁴ Graduate School of Engineering, Faculty of Engineering, Kobe University, Kobe 657-8501, Japan

* Correspondence: chensinan@gold.kobe-u.ac.jp

Abstract

Cross-domain collaborative filtering effectively alleviates the data sparsity issue but raises serious privacy concerns. Federated learning has been integrated into cross-domain collaborative filtering to reduce these risks by securely exchanging embeddings or model parameters. However, the current federated paradigm often relies on the simple alignment of coarse-grained representations, while propagating information on a rigid local graph. Without fine-grained preference modeling, models easily suffer from representation collapse. The inherent sparsity of local graphs further limits their robustness. To address these issues, we propose P2P-IAGR, a privacy-preserving peer-to-peer cross-domain collaborative filtering framework based on intent-adaptive graph reconstruction. Specifically, our framework first disentangles user and item representations into fine-grained latent intent prototypes. We perturb these prototypes using Local Differential Privacy (LDP) and securely exchange them across domains. A contrastive learning strategy is then used for cross-domain alignment. Next, guided by the combined cross-domain intent prior, P2P-IAGR differentially reconstructs an augmented graph view. We apply a dual-view structural contrastive learning objective to dynamically inject external collaborative signals into the sparse local topology. Extensive experiments on real-world datasets show that P2P-IAGR significantly outperforms state-of-the-art methods, achieving an average improvement of 5.85% in NDCG and 6.89% in HR.

Keywords: cross-domain collaborative filtering; federated learning; Local Differential Privacy; intent disentanglement; contrastive learning



Academic Editors: Priyadarsi Nanda and Manoranjan Mohanty

Received: 30 May 2026

Revised: 7 July 2026

Accepted: 14 July 2026

Published: 15 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

In the era of information overload, recommender systems have become essential tools by effectively filtering out irrelevant information for users. While the typical recommendation algorithms (e.g., collaborative filtering) [1,2] have achieved considerable success, they still heavily rely on users' historical interaction records. In real-world applications, users' interaction matrices are often extremely sparse, which makes it difficult for recommender systems to effectively model user preferences. To tackle this challenge, the cross-domain paradigm of collaborative filtering serves as a powerful approach for alleviating data sparsity [3,4]. The main idea is to leverage knowledge sharing and transfer across domains to compensate for the scarcity of data within a single domain [5]. However, these existing methods fundamentally rely on a strong assumption: user data can be fully shared between

the source and target domains. Given the increasingly strict data protection regulations, such a fully shared data interaction paradigm carries a significant risk of leaking sensitive user information.

To mitigate the risk of data leakage, federated learning (FL) has been integrated into cross-domain recommendation (CDR) frameworks, enabling collaborative model training without centralizing raw data. Early federated methods primarily focused on securely exchanging model parameters or perturbed data, such as PriCDR [6] and FedGNN [7]. To alleviate the prohibitive communication costs and gradient-leakage risks, subsequent works like P2FCDR [8] and P²DTR [9] introduced lightweight representation sharing and prototype-based federated learning. Recently, to further enrich user semantics under privacy constraints, some works have incorporated auxiliary information into local modeling. For instance, DT-FedSDC [10] and FUPM [11] leverage review texts to decouple preferences or transfer group prototypes, while FP2CDSR [12] employs temporal modeling for sequential behaviors. Despite the success of these federated methods, a critical challenge remains: How to achieve fine-grained semantic alignment and alleviate data sparsity in privacy-preserving CDR without relying on auxiliary data or central servers?

Specifically, current paradigms still face two limitations. First, coarse-grained preference modeling. While exchanging embeddings or macro-level prototypes effectively reduces communication overhead, current paradigms typically rely on the simple alignment of coarse-grained representations. The lack of explicit disentanglement of fine-grained latent preferences restricts the depth of cross-domain semantic correlations, leaving models highly vulnerable to representation collapse during cross-domain alignment. Second, rigid local topologies. While external knowledge is often incorporated into current federated graph paradigms, information is typically propagated on fixed local interaction graphs. The absence of an adaptive mechanism to dynamically reconstruct inherently sparse local graph structures using cross-domain collaborative signals restricts the ability to learn robust representations.

To bridge these gaps, we propose a privacy-preserving peer-to-peer cross-domain collaborative filtering framework via intent-adaptive graph reconstruction (P2P-IAGR). The core idea of P2P-IAGR is to enable privacy-preserving cross-domain collaboration through intent-level prototype exchange, rather than relying on centralized data aggregation or sharing raw local graph information. Specifically, each domain first employs a simplified graph convolutional network to encode local topological relations and disentangle user and item representations into fine-grained latent intent prototypes. To facilitate secure communication, domains only exchange noisy intent prototypes perturbed via Local Differential Privacy (LDP) to protect raw data. Based on these exchanged prototypes, we further design a cross-domain alignment objective to encourage semantic consistency across domains and construct a shared intent prior. To fully exploit the transferred cross-domain knowledge, we inject the shared intent prior back into local structure modeling. Specifically, guided by the combined intent prior, we differentially reconstruct an augmented graph view using the Gumbel-Sigmoid trick. A dual-view structural contrastive learning objective is introduced as an auxiliary training objective between the original and reconstructed views, which further improves the robustness and consistency of the learned representations.

Our main contributions are summarized as follows:

- We propose P2P-IAGR, a novel privacy-preserving peer-to-peer cross-domain collaborative filtering framework. The core contribution of P2P-IAGR is a unified federated collaborative paradigm that securely exchanges intent-level prototypes to enhance local structure modeling, enabling effective cross-domain recommendation without sharing raw interaction data or model parameters.

- To support reliable prototype transfer, we design an intent disentanglement and cross-domain alignment mechanism, which extracts fine-grained intent prototypes perturbed via Local Differential Privacy (LDP) and aligns them across domains to establish a shared intent prior.
- To better incorporate the transferred cross-domain knowledge into local representation learning, we further introduce an intent-adaptive graph reconstruction strategy together with a structural contrastive learning objective, which enriches sparse local graphs and improves representation robustness.

The rest of this paper is organized as follows: Section 2 reviews the related work. Section 3 formulates the problem and details the proposed P2P-IAGR framework. Section 4 presents the experimental settings, including the datasets, evaluation metrics, and baseline methods. Section 5 provides a detailed analysis of the experimental results. Finally, Section 6 concludes the paper.

2. Related Work

In this section, we review the literature closely related to our research. We first introduce the general development of cross-domain recommendation (CDR) in Section 2.1. Subsequently, in Section 2.2, we discuss existing privacy-preserving CDR methods, highlighting their limitations in representation disentanglement and graph topology. Finally, we summarize the key differences between our proposed P2P-IAGR and representative methods.

2.1. Cross-Domain Recommendation

Cross-domain recommendation is often utilized to alleviate data sparsity and cold-start issues by transferring collaborative knowledge across relevant domains [13,14]. Early studies mainly focused on single-target CDR [15,16], which aims to enrich a sparse target domain using data from a resource-rich source domain. These methods relied on matrix factorization (MF) to capture collaborative filtering signals on overlapping users [17]. Subsequently, inspired by foundational architectures such as NCF [18], deep neural networks were adopted to capture complex non-linear user–item preferences. For instance, DCDCSR [19] bridged different systems by combining MF with a deep mapping network. CoNet [20] introduced cross-connections between feed-forward neural networks to facilitate effective knowledge distillation. More recently, the research shifted toward dual-target CDR [21,22], which is designed to concurrently improve recommendation performance in all participating domains. Notable examples include DDTCDR [23], which designed an attention-based graph framework for two-way preference transfer, and ETL [24], which mapped user representations across domains through equivalent transformations. Graph-based approaches [25,26] also became popular because they directly capture high-order structural links in dual-target scenarios. Despite their success, the aforementioned methods mainly rely on a centralized data-sharing paradigm, assuming that raw user–item interactions across domains can be directly used. This centralized approach creates serious privacy risks.

2.2. Privacy-Preserving Cross-Domain Recommendation

To prevent the direct exchange of raw data, federated learning (FL) [27] and Differential Privacy (DP) have been widely adopted in privacy-preserving recommendation [28–30]. The privacy requirement is particularly important in CDR because the knowledge transfer process increases the risk of data leakage. Early privacy-preserving methods (such as PriCDR [6]) attempted to mitigate privacy leakage by adding noise to the original user–item interaction matrix. Models like FedGNN [7] introduced a federated framework for graph neural networks, sharing gradients protected by Local Differential Privacy (LDP)

to collaboratively train the recommendation model. To avoid the heavy communication costs and gradient-leakage risks, researchers proposed sharing lightweight representations instead of full model parameters. P2FCDR [8] built a peer-to-peer federated network that learns an orthogonal mapping matrix to transform user embeddings securely. Furthermore, prototype-based methods (e.g., FedProto [31]) have gained traction. For instance, P²DTR [9] adopts this paradigm alongside Private Set Intersection (PSI) and Cross-Domain Sharing (CDS) mechanisms for secure knowledge transfer. Recently, some works have incorporated auxiliary information into local interaction modeling. For instance, DT-FedSDC [10] utilizes text semantic enhancement to decouple user preferences into shared and domain-specific parts. Similarly, FUPM [11] integrates auxiliary information such as review texts and leverages federated learning to transfer differentially private user group prototypes across domains. From a temporal perspective, FP2CDSR [12] proposes a centralized federated cross-domain sequential framework that distributes differentially private user representations, employing self-attention and temporal modeling to restore semantic information. In summary, existing privacy-preserving CDR paradigms suffer from coarse-grained representation sharing and rigid local topologies. These limitations hinder explicit intent disentanglement and adaptive graph reconstruction, making models vulnerable to representation collapse and data sparsity. A detailed comparison between our approach and representative existing works is summarized in Table 1.

Table 1. Comparison between our proposed method and existing privacy-preserving cross-domain recommendation methods. (“Aux. Data” refers to texts, reviews, or sequential behaviors. “Fine-grained” denotes explicit fine-grained preference disentanglement. “Dynamic Topo.” denotes dynamic local topology reconstruction. ✓ indicates the feature is utilized, and × indicates it is not.)

Method	Privacy Mechanism	Comm. Paradigm	Aux. Data	Fine-Grained	Dynamic Topo.
PriCDR [6]	Data Perturbation	Interaction Matrix	×	×	×
FedGNN [7]	LDP	Model Gradients	×	×	×
DT-FedSDC [10]	DP	Embeddings	✓ (Texts)	×	×
FUPM [11]	DP	Group Prototypes	✓ (Reviews)	×	×
FP2CDSR [12]	DP	Embeddings	✓ (Sequences)	×	×
P2FCDR [8]	LDP	User Embeddings	×	×	×
P ² DTR [9]	PSI & CDS	Cluster Prototypes	×	×	×
Ours	LDP	Prototypes	× (Pure Graph)	✓	✓

It should be noted that these methods were originally designed with different objectives; the comparison focuses on the dimensions most relevant to our problem setting.

3. Methods

In this section, we detail the proposed privacy-preserving cross-domain recommendation framework, P2P-IAGR. To provide a clear roadmap of our methodology, we first formally define the task and present an overview of the overall architecture. Following this, we elaborate on the core components of our model. Finally, we summarize the optimization objective and provide a theoretical complexity analysis.

3.1. Problem Definition and Overview

We consider a cross-domain recommendation scenario with two distinct domains $\mathcal{D}_A$ and $\mathcal{D}_B$. These domains share a common set of users $\mathcal{U}$, but maintain disjoint item sets $\mathcal{V}_A$ and $\mathcal{V}_B$. Let $\mathbf{R}_A \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{V}_A|}$ and $\mathbf{R}_B \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{V}_B|}$ denote their respective implicit feedback matrices, where an entry equals 1 if an interaction (e.g., click or purchase) exists, and 0 otherwise. Our objective is to predict unobserved entries in both $\mathbf{R}_A$ and $\mathbf{R}_B$ to provide Top- N recommendations. Under strict privacy constraints, raw interaction data must remain localized.

Figure 1 illustrates our P2P-IAGR framework for privacy-preserving decentralized cross-domain recommendation. Taking Domain A as an example, it comprises three primary modules: (1) intent disentanglement, which encodes the local user-item interaction and extracts disentangled intent prototypes S^A to capture fine-grained user preferences. (2) cross-domain contrastive alignment, which aligns local intents with the received LDP-perturbed prototypes $\tilde{S}^B$ through a contrastive alignment mechanism. (3) structure-adaptive contrastive learning, which uses cross-domain intents as a prior to generate an adaptive augmented graph, and performs contrastive learning to inject structural insights. Finally, these three modules are jointly optimized alongside the primary recommendation task.




Figure 1. The overall architecture of the proposed P2P-IAGR framework. It illustrates the workflow within and between two peer clients (Client A and Client B). Each client contains three main components: (1) local structural learning and intent disentanglement, which extracts local intents from the user-item graph via LightGCN; (2) P2P exchange and alignment, where local intents are exchanged under Local Differential Privacy (LDP) noise injection and semantically aligned; and (3) intent-adaptive graph contrastive learning, which utilizes the fused intents to reconstruct graphs and perform contrastive optimization.

3.2. Local Structural Learning and Intent Disentanglement

3.2.1. Local Structural Learning

In our decentralized dual-target cross-domain recommendation (CDR) scenario, Domain A and Domain B are treated as symmetric participants. To capture the collaborative signals within each domain, we employ the simplified graph convolutional architecture, LightGCN [32], as our local representation learning. Let $\mathbf{E}^u \in \mathbb{R}^{|\mathcal{U}| \times D}$ and $\mathbf{E}^v \in \mathbb{R}^{|\mathcal{V}| \times d}$ denote the ID-based embedding matrices for local users $\mathcal{U}$ and items $\mathcal{V}$, where D is the embedding dimension. To efficiently model the high-order connectivity, the local propagation rule at the l -th layer is defined as

$$\begin{aligned} \mathbf{z}_{i,l}^u &= \sum_{j \in \mathcal{N}_i} \frac{1}{\sqrt{|\mathcal{N}_i| |\mathcal{N}_j|}} \mathbf{e}_{j,l-1}^v, \\ \mathbf{z}_{j,l}^v &= \sum_{i \in \mathcal{N}_j} \frac{1}{\sqrt{|\mathcal{N}_j| |\mathcal{N}_i|}} \mathbf{e}_{i,l-1}^u, \end{aligned} \quad (1)$$

where $\mathcal{N}_i$ and $\mathcal{N}_j$ represent the sets of interacted neighbors for user u_i and item v_j in the local domain, respectively. After L layers of propagation, we perform layer aggregation to obtain the refined local structural embeddings $\mathbf{z}_i^u$ and $\mathbf{z}_j^v$:

$$\mathbf{z}_i^u = \sum_{l=0}^L \mathbf{e}_{i,l}^u, \quad \mathbf{z}_j^v = \sum_{l=0}^L \mathbf{e}_{j,l}^v, \quad (2)$$

These embeddings encode the local topological structure. They will be further processed by the subsequent intent-aware disentanglement module to ensure the local model could capture both topological relations and latent behavioral patterns.

3.2.2. Intent Disentanglement

While the local graph convolutional layers capture structural dependencies, they often fail to capture the fine-grained behavioral patterns. To address this, we disentangle the user-item preferences into K latent intents. Especially, we define a set of learnable intent prototypes $\mathbf{S}^u = \{\mathbf{s}_1^u, \dots, \mathbf{s}_K^u\}$ and $\mathbf{S}^v = \{\mathbf{s}_1^v, \dots, \mathbf{s}_K^v\}$ representing K intent clusters for users and items, respectively, where each $\mathbf{s}_k \in \mathbb{R}^D$. These prototypes are shared across domains in the federated network to facilitate global semantic alignment. For a user u_i , we derive the relevance score between their l -th layer structural embedding $\mathbf{z}_{i,l}^u$ and each intent prototype $\mathbf{s}_k^u$ using a softmax-normalized transformation:

$$P(s_k^u | \mathbf{z}_{i,l-1}^u) = \frac{\exp(\mathbf{z}_{i,l-1}^u \top \mathbf{s}_k^u)}{\sum_{k'=1}^K \exp(\mathbf{z}_{i,l-1}^u \top \mathbf{s}_{k'}^u)}. \quad (3)$$

This probability distribution P reflects the degree of a user's preference for different latent intentions. Based on the derived intent distributions, we can generate the intent-aware representation $\mathbf{g}_{i,l}^u$ by aggregating information across all K intents:

$$\mathbf{g}_{i,l}^u = \sum_{k=1}^K \mathbf{s}_k^u P(s_k^u | \mathbf{z}_{i,l}^u). \quad (4)$$

Intuitively, while the structural embeddings capture what items a user interacts with, the intent prototypes aim to discover why they interact. For instance, a user's interactions might be driven by diverse latent motives. By calculating the distribution over K prototypes, Equations (3) and (4) effectively disentangle these mixed motives into distinct, interpretable

preference components. Similarly, the item-side intent representation $\mathbf{g}_{j,l}^v$ is computed following the same logic using $\{\mathbf{s}_k^v\}$. Finally, we refine the node representations by integrating the local collaborative filtering signals with the intent disentangled collaborative relations. The integrated embedding for the l -th layer is formulated as

$$\begin{aligned}\mathbf{e}_{i,l}^u &= \mathbf{z}_{i,l}^u + \mathbf{g}_{i,l}^u, \\ \mathbf{e}_{j,l}^v &= \mathbf{z}_{j,l}^v + \mathbf{g}_{j,l}^v.\end{aligned}\quad (5)$$

By combining the structural message $\mathbf{z}_l$, and the intent-aware message $\mathbf{g}_l$, the framework ensures a robust representation that is resilient to both data sparsity in local domains and potential negative transfer during cross-domain knowledge exchange.

To optimize the recommendation task, we employ the Bayesian Personalized Ranking (BPR) loss. BPR assumes that a user prefers an observed interacted item over an unobserved one, which is formulated as

$$\mathcal{L}_{BPR} = - \sum_{(i,j,j') \in \mathcal{O}} \ln \sigma(\hat{y}_{ij} - \hat{y}_{ij'}) \quad (6)$$

where $\mathcal{O} = \{(i, j, j') \mid \mathbf{R}_{ij} = 1, \mathbf{R}_{ij'} = 0\}$ is the pairwise training set; $\sigma(\cdot)$ is the sigmoid function; and the predicted preference score $\hat{y}_{ij}$ is calculated via the inner product of their final representations: $\hat{y}_{ij} = (\mathbf{e}_i^u)^\top \mathbf{e}_j^v$.

3.3. Decentralized Peer-to-Peer Exchange and Alignment

3.3.1. Privacy Guarantees via Local Differential Privacy

We assume an Honest-but-Curious adversary model, which is standard in federated learning. Specifically, all participating domains honestly follow the cross-domain training protocol, but they may attempt to infer sensitive raw data from the received information. Consequently, the protection scope of our framework includes the raw user–item interaction records and the local graph topological structures. To mitigate potential inference attacks during peer-to-peer communication, our framework prohibits the transmission of raw node embeddings or local graph parameters. Instead, it only exchanges derived intent prototypes after local perturbation.

Specifically, we perturb the transmitted intent prototypes using Local Differential Privacy (LDP). For any local intent prototype $\mathbf{s}_k \in \mathbb{R}^D$ (either from $\mathbf{S}^u$ or $\mathbf{S}^v$), we first bound its L_1 -sensitivity by applying a clipping operation $\bar{\mathbf{s}}_k = \mathbf{s}_k \cdot \min(1, C/\|\mathbf{s}_k\|_1)$, where C is a predefined threshold. This ensures the global L_1 -sensitivity Δf is bounded by $2C$. The perturbed prototype $\tilde{\mathbf{s}}_k$ is then generated as

$$\tilde{\mathbf{s}}_k = \bar{\mathbf{s}}_k + \text{Lap}(0, \lambda)^D \quad (7)$$

where $\text{Lap}(0, \lambda)^D$ denotes a D -dimensional noise vector sampled from the Laplace distribution with scale $\lambda = \Delta f / \epsilon_0$. By setting the noise scale to $\lambda = 2C / \epsilon_0$, each released prototype rigorously satisfies ϵ_0 -LDP. It is worth noting that although perturbed prototypes may still contain high-level aggregated preference information, LDP perturbation prevents reliable inference of any individual user's raw interactions under the defined threat model.

Since the cross-domain semantic alignment requires exchanging the perturbed prototypes $\tilde{\mathbf{s}}_k$ over T communication rounds, the privacy loss accumulates throughout training. According to the sequential composition property of differential privacy, if each communication round satisfies ϵ_0 -LDP, the cumulative privacy budget after T rounds is bounded by

$$\epsilon_{total} \leq T\epsilon_0 \quad (8)$$

While this implies that the overall privacy loss grows linearly with the number of training rounds, our framework benefits from fast convergence. Specifically, the training process typically converges within a small number of communication rounds (e.g., $T \approx 30$). This fast convergence naturally bounds the maximum number of rounds, providing an explicit and controllable upper bound on the cumulative privacy cost without requiring excessive training. In our implementation, we set the per-round privacy budget $\epsilon_0 = 2.0$, which achieves an optimal trade-off between strict privacy guarantees and recommendation utility.

3.3.2. Cross-Domain Semantic Alignment

To ensure that the disentangled intent prototypes capture consistent semantics across heterogeneous domains, we introduce a client-side contrastive alignment scheme. In each communication round, Domain A receives the privacy-preserved intent prototypes $\tilde{\mathbf{S}}^B$ from Domain B. Domain A then incorporates a cross-domain contrastive loss to align its local prototypes $\mathbf{S}^A$ with the remote prototypes $\tilde{\mathbf{S}}^B$. The contrastive objective is formulated using the InfoNCE loss:

$$\mathcal{L}_{CCL}^A = - \sum_{k=1}^K \log \frac{\exp(\text{sim}(\mathbf{s}_k^A, \tilde{\mathbf{s}}_k^B) / \tau)}{\sum_{j=1}^K \exp(\text{sim}(\mathbf{s}_k^A, \tilde{\mathbf{s}}_j^B) / \tau)} \quad (9)$$

where $\text{sim}(\cdot)$ denotes the cosine similarity and τ is the temperature scale. By treating the identically indexed intent $\tilde{\mathbf{s}}_k^B$ as a positive anchor and the other intents $\tilde{\mathbf{s}}_{j \neq k}^B$ as negative samples, this objective forces the local intent prototypes to align with the global semantics. From an intuitive perspective, since Domain A and Domain B possess completely disjoint item sets, their locally learned intent spaces are initially isolated. This contrastive alignment acts as a universal translator, forcing the same prototype index to represent a consistent latent semantic factor across both domains.

It should be noted that we do not assume intent prototypes with identical indices are semantically aligned across domains a priori. The semantic correspondence is not a pre-defined assumption but a learning target progressively established during training. Specifically, driven by this alignment objective and the implicit bridge of overlapping users, the optimization process forces the network to embed shared semantic concepts into the same index to minimize the loss. This alignment effectively mitigates semantic drift during peer-to-peer federated training. Consequently, it ensures that knowledge transferred from Domain B is accurately interpreted by Domain A. Symmetrically, Domain B performs the identical operation to align its local prototypes with the privacy-preserved prototypes $\tilde{\mathbf{S}}^A$ received from Domain A. It optimizes a symmetric contrastive objective, denoted as $\mathcal{L}_{CCL}^B$.

3.3.3. Cross-Domain Intent Fusion

In our decentralized architecture, each domain directly integrates local semantics with privacy-preserved external knowledge. Taking Domain A as an example, we construct a unified cross-domain intent prior to mitigate local data sparsity and inject collaborative signals from Domain B. Specifically, the fused intent prior $\mathbf{C} = \{\mathbf{c}_k\}_{k=1}^K$ is obtained by element-wise adding the local user prototypes $\mathbf{s}_k^A$ and the remote perturbed prototypes $\tilde{\mathbf{s}}_k^B$:

$$\mathbf{c}_k = \mathbf{s}_k^A + \tilde{\mathbf{s}}_k^B \quad (10)$$

3.4. Intent-Adaptive Graph Contrastive Learning

3.4.1. Adaptive Graph Reconstruction

Having the fused cross-domain intent priors $\mathbf{C}$, we proceed to reconstruct the local interaction graph to alleviate local data sparsity. Without loss of generality, we take Domain A as an example. To bridge the inherent semantic gap between user and item embeddings, we project the local node representations into the unified fused intent space. Specifically, we calculate the distribution of local nodes over the cross-domain prior to reconstruct the prior-aware representations:

$$\tilde{\mathbf{e}}_u = \sum_{k=1}^K p_u^k \mathbf{c}_k^u, \quad \tilde{\mathbf{e}}_v = \sum_{k=1}^K q_v^k \mathbf{c}_k^v \quad (11)$$

where p_u^k and q_v^k denote the attention weights of user u and item v toward the k -th fused intent prototype, respectively. These weights are computed via the inner product and the softmax normalization:

$$p_u^k = \frac{\exp((\mathbf{e}_u)^\top \mathbf{c}_k^u)}{\sum_{j=1}^K \exp((\mathbf{e}_u)^\top \mathbf{c}_j^u)}, \quad q_v^k = \frac{\exp((\mathbf{e}_v)^\top \mathbf{c}_k^v)}{\sum_{j=1}^K \exp((\mathbf{e}_v)^\top \mathbf{c}_j^v)} \quad (12)$$

The reconstructed representations $\tilde{\mathbf{e}}_u$ and $\tilde{\mathbf{e}}_v$ map the cross-domain collaborative signals into a unified latent space. Based on these representations, we evaluate the latent connection probability for any user–item pair using the inner product combined with a sigmoid function to generate a continuous probability matrix Ω :

$$\Omega_{u,v} = \sigma\left(\frac{(\tilde{\mathbf{e}}_u)^\top \tilde{\mathbf{e}}_v}{\tau}\right) \quad (13)$$

where τ is a temperature controlling the smoothness of the probability distribution. To incorporate cross-domain latent structures while preserving the original interactions, we perform a soft interpolation between the observed adjacency matrix $\mathbf{A}$ and the inferred probability matrix Ω :

$$\Pi_{u,v} = \lambda \mathbf{A}_{u,v} + (1 - \lambda) \Omega_{u,v} \quad (14)$$

where $\lambda \in [0, 1]$ is the interpolation weight. The resulting matrix Π integrates deterministic local observations with probabilistic cross-domain inferences. To derive a sparse graph while preserving end-to-end differentiability, we apply the Gumbel–Sigmoid trick to sample discrete edges:

$$\hat{\mathbf{A}}_{u,v} = \sigma\left(\frac{\log(\Pi_{u,v}) - \log(1 - \Pi_{u,v}) + g_1 - g_2}{\tau}\right) \quad (15)$$

where $g_1, g_2 \sim \text{Gumbel}(0, 1)$ and τ is the sampling temperature. Consequently, we generate a stochastically perturbed view $\hat{\mathbf{A}}$ for local contrastive learning. The intuitive motivation behind this reconstruction process is to guess and fill in the missing interactions in the sparse local graph using the rich collaborative knowledge transferred from the remote domain. Instead of relying solely on the deterministic but sparse local observations $\mathbf{A}$, we infer potential unobserved edges Ω . The Gumbel–Sigmoid trick then allows us to sample these inferred probabilities into concrete, discrete edges while maintaining the end-to-end differentiability of the neural network. Consequently, we generate a stochastically perturbed view $\hat{\mathbf{A}}$ for local contrastive learning.

3.4.2. Dual-View Graph Contrastive Learning

To effectively transfer cross-domain knowledge from the augmented topology while preserving local collaborative signals, we introduce a dual-view structural contrastive learning (SCL) task. First, we feed the sampled augmented adjacency matrix $\hat{\mathbf{A}}$ into the shared graph encoder. This message-passing process yields the augmented node representations for users and items, denoted as $\hat{\mathbf{e}}_U$ and $\hat{\mathbf{e}}_V$, respectively. Subsequently, we align the augmented representations with the original ones. We treat the representations derived from the explicit original graph $\mathbf{A}$ (i.e., $\mathbf{e}_u$ and $\mathbf{e}_v$) as the anchor views, and the augmented representations (i.e., $\hat{\mathbf{e}}_u$ and $\hat{\mathbf{e}}_v$) as the positive samples. Following the InfoNCE paradigm, other nodes within the current mini-batch $\mathcal{B}$ are treated as negative samples. The contrastive loss for users is formulated as

$$\mathcal{L}_{SCL}^{(u)} = - \sum_{u \in \mathcal{B}} \log \frac{\exp(\cos(\mathbf{e}_u, \hat{\mathbf{e}}_u) / \tau_{scl})}{\sum_{u' \in \mathcal{B}} \exp(\cos(\mathbf{e}_u, \hat{\mathbf{e}}_{u'}) / \tau_{scl})} \quad (16)$$

where $\cos(\cdot, \cdot)$ denotes the cosine similarity function and τ_{scl} is the temperature controlling the penalty strength on hard negative samples. Symmetrically, the contrastive loss for items, denoted as $\mathcal{L}_{SCL}^{(v)}$, is computed in the same manner. The overall structure-adaptive contrastive loss is defined as the sum of both sides:

$$\mathcal{L}_{SCL} = \mathcal{L}_{SCL}^{(u)} + \mathcal{L}_{SCL}^{(v)} \quad (17)$$

By optimizing $\mathcal{L}_{SCL}$, the model maximizes the mutual information between the localized collaborative signals and the cross-domain augmented views.

3.5. Model Learning and Federated Optimization

The final local objective function for each client (e.g., Domain A) is a joint optimization of the primary recommendation task, the cross-domain intent alignment, and the dual-view structural contrastive learning:

$$\mathcal{L}_{final} = \mathcal{L}_{BPR} + \beta \mathcal{L}_{CCL} + \gamma \mathcal{L}_{SCL} \quad (18)$$

where β and γ are hyper-parameters controlling the strength of semantic alignment and structural self-supervision, respectively.

To optimize this objective within our decentralized peer-to-peer (P2P) architecture, Domain A and B collaboratively train their models without sharing raw interaction data. As summarized in Algorithm 1, the overall training procedure alternates between cross-domain communication and local optimization. Specifically, in each communication round, both domains independently extract their local intent prototypes and apply the Local Differential Privacy (LDP) mechanism to generate the perturbed prototypes $\tilde{\mathbf{S}}^A$ and $\tilde{\mathbf{S}}^B$. After exchanging these noisy prototypes, each domain constructs the fused intent prior, derives the augmented graph view, and performs local backpropagation to minimize its respective joint objective ($\mathcal{L}_{final}^A$ or $\mathcal{L}_{final}^B$). This iterative process enables both models to capture cross-domain insights while strictly preserving privacy.

Algorithm 1 Optimization algorithm for P2P-IAGR

Require: Local graphs $\mathcal{G}^A, \mathcal{G}^B$; communication rounds R ; local epochs E ; learning rate η ; privacy budget ϵ ; number of intents K ; loss weights β, γ .

Ensure: Optimized local model parameters Θ^A and Θ^B .

- 1: Initialize parameters Θ^A and Θ^B for Domain A and B.
- 2: **for** round $r = 1, 2, \dots, R$ **do**
- 3: **Phase 1: Prototype Extraction & LDP Perturbation**
- 4: Extract clean prototypes $\mathbf{S}^A$ and $\mathbf{S}^B$.
- 5: $\tilde{\mathbf{S}}^* \leftarrow \mathbf{S}^* \cdot \min(1, C/\|\mathbf{S}^*\|_1)$ for $* \in \{A, B\}$.
- 6: $\tilde{\mathbf{S}}^A \leftarrow \tilde{\mathbf{S}}^A + \text{Lap}(0, 2C/\epsilon)^D$
- 7: $\tilde{\mathbf{S}}^B \leftarrow \tilde{\mathbf{S}}^B + \text{Lap}(0, 2C/\epsilon)^D$
- 8: **Phase 2: Peer-to-Peer Exchange**
- 9: Domain A sends $\tilde{\mathbf{S}}^A$ to B; Domain B sends $\tilde{\mathbf{S}}^B$ to A.
- 10: **Phase 3: Local Federated Optimization**
- 11: **for** epoch $e = 1, 2, \dots, E$ on Domain A (Parallel on B) **do**
- 12: **Step 1: Structural Learning**
- 13: Obtain structural embeddings $\mathbf{Z}^u, \mathbf{Z}^v$ via LightGCN.
- 14: Update node embeddings: $\mathbf{e}_i \leftarrow \mathbf{z}_i + \mathbf{g}_i$.
- 15: **Step 2: Semantic Alignment**
- 16: Calculate $\mathcal{L}_{CCL}^A$ aligning local $\mathbf{S}^A$ with remote $\tilde{\mathbf{S}}^B$.
- 17: **Step 3: Graph Reconstruction**
- 18: Compute fused intent prior $\mathbf{C} \leftarrow \mathbf{S}^A + \tilde{\mathbf{S}}^B$.
- 19: Infer probability Ω^A and interpolate Π^A .
- 20: Sample augmented graph $\hat{\mathbf{A}}^A$ via Gumbel–Sigmoid.
- 21: **Step 4: Contrastive Learning**
- 22: Compute structure-adaptive contrastive loss $\mathcal{L}_{SCL}^A$.
- 23: **Step 5: Joint Update**
- 24: $\mathcal{L}_{final}^A \leftarrow \mathcal{L}_{BPR}^A + \beta \mathcal{L}_{CCL}^A + \gamma \mathcal{L}_{SCL}^A$.
- 25: Update parameters: $\Theta^A \leftarrow \Theta^A - \eta \nabla \mathcal{L}_{final}^A$.
- 26: **end for**
- 27: **end for**
- 28: **return** Θ^A, Θ^B

3.6. Complexity Analysis

In this section, we analyze the time and space complexity of P2P-IAGR under one local domain, since the two domains are trained symmetrically and thus have the same order of complexity. Let $|\mathcal{U}|$, $|\mathcal{V}|$, and $|\mathcal{E}|$ denote the numbers of users, items, and observed interactions in the local domain, respectively. Let D be the embedding dimension, L the number of graph propagation layers, K the number of disentangled intent prototypes, $|\mathcal{B}|$ the mini-batch size, and T the number of communication rounds.

3.6.1. Time Complexity

The time cost mainly consists of three parts: local structural learning, decentralized communication, and graph reconstruction. For local structural learning, the LightGCN encoding and intent disentanglement take $\mathcal{O}(L|\mathcal{E}|D + L(|\mathcal{U}| + |\mathcal{V}|)KD)$. In the decentralized stage, transmitting and aligning K perturbed intent prototypes require $\mathcal{O}(KD + K^2D)$ per round, which is significantly more efficient than exchanging full node embeddings. For adaptive graph reconstruction, evaluating all user–item connections would incur a worst-case cost of $\mathcal{O}(|\mathcal{U}||\mathcal{V}|D)$. However, restricting this step to sampled candidates or sparse neighborhoods significantly reduces the actual computational overhead. Subsequently, dual-view graph contrastive learning introduces $\mathcal{O}(L|\hat{\mathcal{E}}|D + |\mathcal{B}|^2D)$ for encoding the reconstructed graph (with $|\hat{\mathcal{E}}|$ edges) and mini-batch similarity computation. Overall, considering T communication rounds, the total time complexity is $\mathcal{O}(T \cdot (L(|\mathcal{E}| + |\hat{\mathcal{E}}|)D + L(|\mathcal{U}| + |\mathcal{V}|)KD + K^2D + |\mathcal{B}|^2D))$.

3.6.2. Space Complexity

The space complexity is primarily determined by the storage of embeddings, prototypes, graph structures, and intermediate representations. Specifically, storing user and item embeddings requires $\mathcal{O}((|\mathcal{U}| + |\mathcal{V}|)D)$, while the intent prototypes require only a negligible $\mathcal{O}(KD)$. Under sparse representation, the original local interaction graph and the reconstructed graph require $\mathcal{O}(|\mathcal{E}|)$ and $\mathcal{O}(|\hat{\mathcal{E}}|)$ memory, respectively. In addition, maintaining the intermediate node embeddings across L layers for backpropagation incurs $\mathcal{O}(L(|\mathcal{U}| + |\mathcal{V}|)D)$. Therefore, the overall space complexity is dominated by the node embedding storage and sparse graph structures, making the proposed framework highly memory-efficient for decentralized recommendation on large-scale sparse graphs.

4. Experimental Settings

In this section, we first introduce the real-world datasets utilized in our evaluation. Next, we present the state-of-the-art baseline methods for comparison. We then detail the evaluation protocol and metrics used to assess the model performance. Finally, we provide the implementation details and hyper-parameter settings to ensure the reproducibility of our experiments.

4.1. Datasets

We evaluate our proposed method on six representative domains from the real-world Amazon dataset: Cell Phones and Accessories (Cell), Electronics (Elec), Movies and TV (Movie), Video Games (Game), Grocery and Gourmet Food (Food), and Home and Kitchen (Kitchen). Following previous work [8], we construct three cross-domain pairs: Cell–Elec, Movie–Game, and Food–Kitchen. For each dataset, we retain overlapping users who have historical interactions in both the source and target domains. During data preprocessing, all explicit ratings are converted into implicit feedback, where an entry is set to one if a user has interacted with an item, and zero otherwise. Consistent with prior studies [33], we further filter the datasets by removing users with fewer than five interactions and items with fewer than 10 interactions. The detailed statistics of the datasets are summarized in Table 2.

Table 2. The statistics of datasets.

Domain Pair	Domain	Shared Users	Items	Ratings	Sparsity
Elec & Cell	Elec	14,591	11,270	397,034	99.76%
	Cell		19,977	428,493	99.85%
Game & Movie	Game	8252	8737	196,792	99.73%
	Movie		9218	284,589	99.63%
Food & Kitchen	Food	2424	3251	57,404	99.27%
	Kitchen		2880	53,498	99.23%

4.2. Baselines

To evaluate the effectiveness of our proposed method, we compare it against several representative state-of-the-art baselines, which can be grouped into three groups.

Single-Domain Recommendation:

- NeuMF [18]: A widely used neural collaborative filtering framework that replaces the conventional inner product with a multi-layer perceptron.
- LightGCN [32]: A light graph convolution model that eliminates feature transformation and non-linear activation.

Cross-Domain Recommendation:

- CoNet [20]: A collaborative cross-network architecture that facilitates bidirectional knowledge transfer through cross-connections between feed-forward networks.
 - DDTCDR [22]: A deep dual transfer model that leverages orthogonal mappings to transfer user preferences and structural knowledge across different domains.
 - ETL [24]: An advanced CDR approach that employs an equivalent transformation learning framework to model the joint distribution of cross-domain user behaviors.
- Privacy-Preserving Recommendation:
- FedGNN [7]: A federated graph neural network framework designed for privacy-preserving recommendation.
 - PriCDR [6]: A privacy-preserving CDR model that shares a differentially private rating matrix from the source domain to the target domain.
 - P2FCDR [8]: A privacy-preserving federated cross-domain recommendation framework that learns cross-domain embedding transformations via an optimizable orthogonal mapping matrix.
 - P²DTR [9]: A privacy-preserving dual-target cross-domain recommendation framework that utilizes the Private Set Intersection (PSI) algorithm and prototype-based federated learning to collaboratively model shared knowledge.

4.3. Evaluation Protocol and Metrics

For evaluation, we adopt the widely used leave-one-out protocol [18]. Specifically, for each user, we select the most recent interaction as the test set, the second most recent interaction as the validation set for hyper-parameter tuning, and the remaining historical interactions as the training set. Furthermore, during the evaluation phase, we pair each ground-truth item in the test set with 99 randomly sampled unobserved items as negative instances. Our main goal is to evaluate the performance of Top- K recommendation in the privacy-preserving scenario. To this end, we employ three standard evaluation metrics: Hit Ratio (HR@ K), Normalized Discounted Cumulative Gain (NDCG@ K), and Mean Reciprocal Rank (MRR). Among these, HR@ K measures whether the true item is successfully retrieved within the top K positions of the ranked list. NDCG@ K and MRR are used to evaluate the ranking of the hit items. Naturally, higher scores in these metrics correspond to more accurate recommendation capabilities.

4.4. Hyper-Parameter Settings

For equitable evaluation across all baselines, we adopted a standardized configuration. Specifically, the latent vector size D and mini-batch size were fixed at 64 and 1024, respectively. Optimization was carried out using the Adam optimizer with a weight decay of 0.001. For all GNN-based methods, the propagation depth L was restricted to three. Within our specific modules, the representation space was partitioned into $K = 16$ latent intents. We set the regularization coefficients for both cross-domain (β) and structural (γ) contrastive learning to 0.1. We specified an LDP noise scale (ϵ) of 2.0 alongside an alignment temperature (τ) of 0.1. In the process of structure-adaptive augmentation, the soft edge probability parameter (λ) is set to 0.5 and the structural scaling factor (τ_{scl}) is set to 0.15.

5. Results and Discussion

In this section, we conduct extensive experiments on real-world datasets to evaluate the effectiveness of our proposed framework. Specifically, we aim to answer the following Research Questions (RQs):

RQ1: How does our proposed framework perform compared to state-of-the-art baselines in cross-domain recommendation tasks?

RQ2: How do the core components (i.e., intent disentanglement, cross-domain contrastive learning, and structure-adaptive graph contrastive learning) contribute to the overall performance?

RQ3: Can our model effectively balance the trade-off between strict privacy protection and recommendation accuracy?

RQ4: How does our framework perform on highly sparse data compared to existing methods?

RQ5: How sensitive is our model to key hyper-parameters?

5.1. Overall Performance Comparisons (RQ1)

To answer RQ1, we present the overall performance of all compared methods in Table 3. Specifically, the table reports the average performance and standard deviation over five independent runs. To ensure the reliability of our results, we conduct a paired *t*-test between our model and the best-performing baseline. The significance markers (* for $p < 0.01$) demonstrate that the performance improvements achieved by our method are statistically significant and not due to random variance. Based on these results, we can make the following observations. Our proposed model P2P-IAGR outperforms all baselines across all evaluation metrics and datasets, which effectively answers RQ1. Specifically, our proposed method achieves an average improvement of 6.5% across all scenarios. For instance, in the Kitchen-domain recommendation, our model achieves 21.92% improvement in *HR@1* and 9.73% in *HR@10*. Similar improvements are also observed in the Movie domain (e.g., 10.48% in *NDCG@10*). These results demonstrate the superior effectiveness of our approach in generating accurate Top-*K* recommendations. Furthermore, cross-domain recommendation models generally outperform single-domain models (e.g., NeuMF and LightGCN). This suggests that leveraging collaborative knowledge from other domains can effectively alleviate the data sparsity issue. Secondly, the privacy-preserving baselines (i.e., FedGNN, P²DTR, and P2FCDR) generally achieve better competitive performance compared to the traditional centralized cross-domain baselines (e.g., ETL and DDTCDR). This demonstrates that decentralized architectures do not compromise the recommendation effectiveness. By integrating advanced graph modeling into the federated paradigm, these frameworks can effectively capture high-order collaborative signals and complex user-item interactions. Finally, among all privacy-preserving baselines, our proposed framework consistently achieves the best performance. This superiority can be attributed to P2P-IAGR's ability to capture fine-grained user behaviors by disentangling preferences into latent intent prototypes. Furthermore, Local Differential Privacy (LDP) noise effectively serves as stochastic data augmentation to align prototypes and prevent semantic drift. Building on these cross-domain prior, our structure-adaptive graph contrastive learning further leverages global semantics to enhance local topologies, thereby alleviating data sparsity and achieves an optimal privacy-utility trade-off.

Table 3. Performance comparison (%) across three bidirectional federated DTCDR scenarios: Food–Kitchen, Game–Movie, and Cell–Elec. The highest value is highlighted in bold, while the second-highest is underlined. Values are reported as mean $\pm$ standard deviation over five independent runs.

Methods	Food-domain recommendation						Kitchen-domain recommendation					
	MRR	NDCG		HR			MRR	NDCG		HR		
		@5	@10	@1	@5	@10		@5	@10	@1	@5	@10
NeuMF	6.12 $\pm$ 0.35	5.17 $\pm$ 0.22	6.42 $\pm$ 0.31	2.13 $\pm$ 0.18	7.45 $\pm$ 0.45	11.82 $\pm$ 0.52	2.84 $\pm$ 0.12	1.93 $\pm$ 0.28	3.32 $\pm$ 0.35	0.87 $\pm$ 0.14	4.15 $\pm$ 0.41	7.38 $\pm$ 0.55
LightGCN	6.41 $\pm$ 0.09	5.38 $\pm$ 0.12	6.75 $\pm$ 0.15	2.44 $\pm$ 0.08	7.82 $\pm$ 0.18	12.43 $\pm$ 0.20	3.22 $\pm$ 0.11	2.18 $\pm$ 0.14	3.61 $\pm$ 0.17	0.96 $\pm$ 0.09	4.62 $\pm$ 0.16	7.54 $\pm$ 0.19
CoNet	6.59 $\pm$ 0.45	5.65 $\pm$ 0.32	7.18 $\pm$ 0.31	2.65 $\pm$ 0.25	8.27 $\pm$ 0.38	13.12 $\pm$ 0.56	3.55 $\pm$ 0.22	2.35 $\pm$ 0.39	3.89 $\pm$ 0.35	1.07 $\pm$ 0.48	4.98 $\pm$ 0.31	7.76 $\pm$ 0.28
DDTCDR	6.78 $\pm$ 0.28	5.85 $\pm$ 0.35	7.29 $\pm$ 0.22	2.78 $\pm$ 0.29	8.42 $\pm$ 0.42	13.55 $\pm$ 0.29	3.62 $\pm$ 0.36	2.58 $\pm$ 0.21	4.05 $\pm$ 0.38	1.10 $\pm$ 0.21	5.16 $\pm$ 0.34	8.09 $\pm$ 0.30
ETL	6.85 $\pm$ 0.31	5.92 $\pm$ 0.38	7.58 $\pm$ 0.45	2.89 $\pm$ 0.22	8.71 $\pm$ 0.55	13.82 $\pm$ 0.46	3.95 $\pm$ 0.29	2.85 $\pm$ 0.41	4.14 $\pm$ 0.31	1.27 $\pm$ 0.24	5.08 $\pm$ 0.37	8.43 $\pm$ 0.59
FedGNN	6.86 $\pm$ 0.24	5.93 $\pm$ 0.31	7.69 $\pm$ 0.29	2.63 $\pm$ 0.27	8.82 $\pm$ 0.36	14.10 $\pm$ 0.33	4.05 $\pm$ 0.21	3.02 $\pm$ 0.36	4.25 $\pm$ 0.23	1.34 $\pm$ 0.36	5.32 $\pm$ 0.29	8.88 $\pm$ 0.26
PriCDR	7.05 $\pm$ 0.57	6.03 $\pm$ 0.34	7.82 $\pm$ 0.41	2.88 $\pm$ 0.32	8.63 $\pm$ 0.49	14.51 $\pm$ 0.57	3.84 $\pm$ 0.25	3.11 $\pm$ 0.39	4.29 $\pm$ 0.46	1.46 $\pm$ 0.19	5.43 $\pm$ 0.52	9.04 $\pm$ 0.39
P ² DTR	7.15 $\pm$ 0.22	6.05 $\pm$ 0.29	7.93 $\pm$ 0.36	3.02 $\pm$ 0.24	8.92 $\pm$ 0.23	14.32 $\pm$ 0.35	3.92 $\pm$ 0.19	3.28 $\pm$ 0.33	4.32 $\pm$ 0.34	1.37 $\pm$ 0.23	5.63 $\pm$ 0.36	8.83 $\pm$ 0.23
P2FCDR	7.08 $\pm$ 0.26	<u>6.15</u> $\pm$ 0.33	7.68 $\pm$ 0.40	<u>3.12</u> $\pm$ 0.18	8.67 $\pm$ 0.27	14.25 $\pm$ 0.34	3.86 $\pm$ 0.23	3.23 $\pm$ 0.27	<u>4.42</u> $\pm$ 0.34	1.28 $\pm$ 0.37	5.55 $\pm$ 0.15	8.72 $\pm$ 0.27
Ours	7.38 $\pm$ 0.18 *	6.48 $\pm$ 0.27 *	8.21 $\pm$ 0.34 *	3.35 $\pm$ 0.12 *	9.18 $\pm$ 0.34 *	15.65 $\pm$ 0.28 *	4.28 $\pm$ 0.18 *	3.49 $\pm$ 0.31 *	4.88 $\pm$ 0.38 *	1.78 $\pm$ 0.14 *	5.89 $\pm$ 0.24 *	9.92 $\pm$ 0.21 *
Improv.	3.22%	5.37%	3.53%	7.37%	2.91%	7.86%	5.68%	6.40%	10.41%	21.92%	4.62%	9.73%
Methods	Game-domain recommendation						Movie-domain recommendation					
	MRR	NDCG		HR			MRR	NDCG		HR		
		@5	@10	@1	@5	@10		@5	@10	@1	@5	@10
NeuMF	4.88 $\pm$ 0.29	4.15 $\pm$ 0.24	5.45 $\pm$ 0.33	1.95 $\pm$ 0.19	6.17 $\pm$ 0.47	10.54 $\pm$ 0.54	3.14 $\pm$ 0.14	1.92 $\pm$ 0.30	3.39 $\pm$ 0.37	1.05 $\pm$ 0.16	3.99 $\pm$ 0.43	6.64 $\pm$ 0.57
LightGCN	5.29 $\pm$ 0.11	4.41 $\pm$ 0.13	5.82 $\pm$ 0.16	2.18 $\pm$ 0.09	6.60 $\pm$ 0.19	10.94 $\pm$ 0.18	3.35 $\pm$ 0.12	2.21 $\pm$ 0.15	3.69 $\pm$ 0.18	1.14 $\pm$ 0.10	4.50 $\pm$ 0.17	7.15 $\pm$ 0.20
CoNet	5.55 $\pm$ 0.38	4.74 $\pm$ 0.31	6.11 $\pm$ 0.26	2.21 $\pm$ 0.39	7.11 $\pm$ 0.45	11.25 $\pm$ 0.32	3.79 $\pm$ 0.42	2.39 $\pm$ 0.28	4.11 $\pm$ 0.39	1.39 $\pm$ 0.22	4.86 $\pm$ 0.29	7.77 $\pm$ 0.46
DDTCDR	5.92 $\pm$ 0.28	4.93 $\pm$ 0.31	6.50 $\pm$ 0.35	2.40 $\pm$ 0.22	7.38 $\pm$ 0.38	11.70 $\pm$ 0.42	3.91 $\pm$ 0.25	2.86 $\pm$ 0.29	4.33 $\pm$ 0.33	1.45 $\pm$ 0.21	5.29 $\pm$ 0.36	8.26 $\pm$ 0.44
ETL	6.10 $\pm$ 0.24	5.27 $\pm$ 0.42	6.75 $\pm$ 0.51	2.46 $\pm$ 0.26	7.82 $\pm$ 0.21	12.31 $\pm$ 0.24	4.15 $\pm$ 0.29	3.12 $\pm$ 0.42	4.66 $\pm$ 0.35	1.59 $\pm$ 0.48	5.65 $\pm$ 0.43	8.90 $\pm$ 0.29
FedGNN	6.45 $\pm$ 0.32	5.35 $\pm$ 0.25	6.88 $\pm$ 0.42	2.68 $\pm$ 0.31	8.09 $\pm$ 0.35	12.98 $\pm$ 0.44	4.33 $\pm$ 0.34	3.65 $\pm$ 0.29	4.78 $\pm$ 0.21	1.73 $\pm$ 0.33	6.06 $\pm$ 0.26	9.42 $\pm$ 0.39
PriCDR	6.24 $\pm$ 0.45	4.62 $\pm$ 0.28	6.76 $\pm$ 0.31	2.48 $\pm$ 0.35	7.57 $\pm$ 0.22	11.09 $\pm$ 0.41	4.44 $\pm$ 0.31	3.04 $\pm$ 0.49	4.47 $\pm$ 0.56	1.71 $\pm$ 0.46	5.18 $\pm$ 0.35	8.61 $\pm$ 0.38
P ² DTR	6.63 $\pm$ 0.28	5.42 $\pm$ 0.26	7.33 $\pm$ 0.21	2.52 $\pm$ 0.28	8.43 $\pm$ 0.34	12.82 $\pm$ 0.38	4.53 $\pm$ 0.25	3.62 $\pm$ 0.37	5.03 $\pm$ 0.29	1.78 $\pm$ 0.38	6.13 $\pm$ 0.42	9.27 $\pm$ 0.26
P2FCDR	6.53 $\pm$ 0.25	<u>5.53</u> $\pm$ 0.37	7.23 $\pm$ 0.22	2.64 $\pm$ 0.31	8.24 $\pm$ 0.23	12.81 $\pm$ 0.27	4.48 $\pm$ 0.26	<u>3.72</u> $\pm$ 0.18	4.91 $\pm$ 0.30	<u>1.83</u> $\pm$ 0.21	5.88 $\pm$ 0.24	9.21 $\pm$ 0.35
Ours	6.87 $\pm$ 0.21 *	5.96 $\pm$ 0.19 *	7.52 $\pm$ 0.26 *	2.83 $\pm$ 0.14 *	9.08 $\pm$ 0.22 *	13.51 $\pm$ 0.15 *	4.77 $\pm$ 0.22 *	4.11 $\pm$ 0.23 *	5.21 $\pm$ 0.34 *	1.98 $\pm$ 0.23 *	6.28 $\pm$ 0.20 *	9.97 $\pm$ 0.26 *
Improv.	3.62%	7.78%	2.59%	5.60%	7.71%	4.08%	5.30%	10.48%	3.58%	8.20%	2.45%	5.84%
Methods	Elec-domain recommendation						Cell-domain recommendation					
	MRR	NDCG		HR			MRR	NDCG		HR		
		@5	@10	@1	@5	@10		@5	@10	@1	@5	@10
NeuMF	2.05 $\pm$ 0.34	1.39 $\pm$ 0.31	1.97 $\pm$ 0.29	0.41 $\pm$ 0.27	2.99 $\pm$ 0.23	4.04 $\pm$ 0.45	7.09 $\pm$ 0.23	5.99 $\pm$ 0.27	7.42 $\pm$ 0.34	3.13 $\pm$ 0.25	8.39 $\pm$ 0.40	12.48 $\pm$ 0.33
LightGCN	2.28 $\pm$ 0.08	1.68 $\pm$ 0.11	2.22 $\pm$ 0.14	0.45 $\pm$ 0.08	3.22 $\pm$ 0.17	4.25 $\pm$ 0.19	7.61 $\pm$ 0.10	6.62 $\pm$ 0.13	7.91 $\pm$ 0.16	3.31 $\pm$ 0.09	8.94 $\pm$ 0.15	13.19 $\pm$ 0.18
CoNet	2.48 $\pm$ 0.38	1.90 $\pm$ 0.33	2.47 $\pm$ 0.26	0.58 $\pm$ 0.38	3.40 $\pm$ 0.21	4.55 $\pm$ 0.37	8.10 $\pm$ 0.54	6.92 $\pm$ 0.41	8.21 $\pm$ 0.25	3.50 $\pm$ 0.43	9.07 $\pm$ 0.58	14.24 $\pm$ 0.41
DDTCDR	2.54 $\pm$ 0.24	2.10 $\pm$ 0.22	2.56 $\pm$ 0.35	0.69 $\pm$ 0.27	3.52 $\pm$ 0.19	4.64 $\pm$ 0.22	8.52 $\pm$ 0.29	7.11 $\pm$ 0.36	8.62 $\pm$ 0.31	3.60 $\pm$ 0.39	10.61 $\pm$ 0.23	14.62 $\pm$ 0.25
ETL	2.88 $\pm$ 0.42	2.32 $\pm$ 0.55	2.79 $\pm$ 0.29	0.84 $\pm$ 0.41	3.68 $\pm$ 0.25	4.82 $\pm$ 0.49	8.62 $\pm$ 0.56	7.39 $\pm$ 0.33	8.84 $\pm$ 0.28	3.91 $\pm$ 0.26	11.79 $\pm$ 0.60	15.94 $\pm$ 0.32
FedGNN	2.98 $\pm$ 0.29	2.44 $\pm$ 0.31	2.95 $\pm$ 0.28	1.27 $\pm$ 0.36	3.68 $\pm$ 0.25	5.04 $\pm$ 0.32	8.41 $\pm$ 0.42	8.18 $\pm$ 0.25	9.02 $\pm$ 0.32	4.18 $\pm$ 0.25	12.09 $\pm$ 0.38	16.50 $\pm$ 0.35
PriCDR	2.95 $\pm$ 0.46	2.34 $\pm$ 0.33	2.92 $\pm$ 0.40	0.92 $\pm$ 0.39	3.40 $\pm$ 0.28	5.10 $\pm$ 0.46	7.95 $\pm$ 0.34	7.42 $\pm$ 0.28	8.94 $\pm$ 0.45	3.89 $\pm$ 0.28	11.58 $\pm$ 0.51	16.13 $\pm$ 0.58
P ² DTR	3.13 $\pm$ 0.21	2.45 $\pm$ 0.28	3.08 $\pm$ 0.35	1.17 $\pm$ 0.23	3.47 $\pm$ 0.42	5.23 $\pm$ 0.29	8.68 $\pm$ 0.28	8.17 $\pm$ 0.32	9.43 $\pm$ 0.39	4.07 $\pm$ 0.22	12.53 $\pm$ 0.25	16.87 $\pm$ 0.22
P2FCDR	3.03 $\pm$ 0.25	<u>2.55</u> $\pm$ 0.22	2.98 $\pm$ 0.29	1.18 $\pm$ 0.37	3.42 $\pm$ 0.26	5.09 $\pm$ 0.43	8.55 $\pm$ 0.32	<u>8.28</u> $\pm$ 0.26	9.23 $\pm$ 0.23	4.02 $\pm$ 0.36	12.39 $\pm$ 0.19	<u>17.09</u> $\pm$ 0.26
Ours	3.22 $\pm$ 0.20 *	2.68 $\pm$ 0.26 *	3.23 $\pm$ 0.33 *	1.35 $\pm$ 0.17 *	3.88 $\pm$ 0.19 *	5.54 $\pm$ 0.27 *	8.95 $\pm$ 0.19 *	8.48 $\pm$ 0.23 *	10.15 $\pm$ 0.27 *	4.38 $\pm$ 0.21 *	13.18 $\pm$ 0.18 *	18.47 $\pm$ 0.22 *
Improv.	2.88%	5.10%	4.87%	6.30%	5.43%	5.93%	3.11%	2.42%	7.64%	4.78%	5.19%	8.07%

* indicates statistical significance at $p < 0.01$ compared to the best baseline, evaluated via paired t -test.

5.2. Ablation Study (RQ2)

To investigate the contribution of each proposed module (RQ2), we design three variants of our model:

- w/o Intent: Replaces the intent disentanglement module with K randomly initialized, learnable vectors, while keeping all subsequent modules unchanged.
- w/o CCL: Removes the cross-domain contrastive learning (CCL) alignment and replaces it with an alignment loss based on l2 distance.
- w/o SCL: Removes structure-adaptive graph contrastive learning, while the rest of the architecture remains unchanged.

The results are shown in Figure 2. We observe that removing any component leads to a noticeable performance degradation. Specifically, the w/o SCL variant suffers the greatest performance degradation across all metrics. For example, removing SCL causes a significant drop in the Game domain, where $NDCG@10$ drops from 7.49% to 6.96%. This indicates that explicitly injecting cross-domain knowledge into the local topological structure through SCL module is crucial for alleviating local data sparsity. By reconstructing the global knowledge to restructure the local model topology, local models can capture high-order collaborative signals. The w/o intent variant exhibits the second largest performance decline. In the Kitchen domain, its MRR drops from 4.26% to 3.82%, and in the Movie domain, $HR@10$ decreases from 9.94% to 9.50%. This confirms that simply using randomly

initialized vectors fails to capture the complex, underlying semantic structures of user behaviors. Without intent disentanglement, the model would learn entangled, coarse-grained representations. Therefore, extracting fine-grained, transferable intent prototypes is indispensable for accurately capturing user preferences and effectively guiding the subsequent adaptive graph reconstruction. Although the performance decline of w/o CCL is relatively the smallest, it still consistently under-performs the full model. This indicates that cross-domain contrastive alignment is important for pulling noisy representations back into the correct semantic space and bridging semantic gaps across different domains. The full model P2P-IAGR consistently achieves the best performance across all datasets and evaluation metrics. This confirms that the intent disentanglement, semantic alignment, and structural graph reconstruction are highly complementary. They work together to address the core challenges of decentralized cross-domain recommendation: alleviating data sparsity, bridging semantic gaps, and preserving high-order structural knowledge.




Figure 2. Performance comparison of P2P-IAGR and its variants on two datasets. The results illustrate the individual contributions of the intent disentanglement (w/o intent), CCL, and SCL modules across MRR, HR@10, and NDCG@10 metrics.

5.3. Privacy–Utility Trade-Off (RQ3)

In FedCDR, there is an inherent trade-off between privacy protection and recommendation accuracy. To answer RQ3, we adjust the per-round privacy budget ϵ_0 in the Local Differential Privacy (LDP) mechanism to investigate this trade-off, while keeping the clipping threshold fixed (e.g., $C = 1.0$). A smaller ϵ_0 corresponds to stronger perturbation, and thus stronger privacy protection in each communication round. Meanwhile, under communication over T rounds, the cumulative privacy budget grows according to the sequential composition property, i.e., $\epsilon_{total} \leq T\epsilon_0$. Therefore, adjusting ϵ_0 directly reflects the trade-off between the overall privacy cost and the recommendation utility throughout training. To study the impact of ϵ_0 , we evaluate the model with $\epsilon_0 \in \{0.1, 0.5, 1.0, 2.0, 5.0\}$. As shown in Figure 3, the recommendation performance consistently improves as ϵ_0 increases from 0.1 to 2.0. This is expected because a larger per-round privacy budget introduces less perturbation into the exchanged intent prototypes, thereby preserving more useful cross-domain semantic information. In particular, when ϵ_0 is at a minimal value (e.g., 0.1), the noise disrupts the feature distribution of the original data, causing the lowest performance. Conversely, the performance of the model improves significantly and reaches its best when ϵ_0 is 2.0. Interestingly, when ϵ_0 further increases to 5.0, the performance exhibits a slight decline. This may be attributed to the fact that the model developed a

certain degree of regularization dependence on specific scales of noise during the training process. Insufficient noise may weaken its generalization ability. Since a smaller ϵ_0 provides stronger privacy guarantees and prevents the accumulated total budget ϵ_{total} from growing excessively over multiple rounds, we need to carefully balance recommendation performance and privacy protection. Thus, we prefer to set the per-round budget ϵ_0 to 2.0. Given our framework's fast convergence (typically within $T \approx 30$ rounds), this setting provides a practical balance between recommendation utility and cumulative privacy cost under our experimental configuration.




Figure 3. Impact of privacy budget ϵ_0 on recommendation performance across three cross-domain tasks. We vary ϵ_0 from 0.1 to 5.0. A smaller ϵ_0 indicates stricter privacy protection but generally leads to a decrease in recommendation accuracy, illustrating the inherent privacy–utility trade-off.

5.4. Alleviating Data Sparsity (RQ4)

To verify whether our model has alleviated the data sparsity issue, we conducted an analysis by changing the proportion of the training data. Specifically, for both domains in the dual-target scenario, we randomly sample 20%, 40%, 60%, and 80% of the original training interactions to construct graph topological structures with different degrees of sparsity. Figure 4 shows the average performance comparisons in terms of MRR, HR@10, and NDCG@10 on the Food–Kitchen and Game–Movie datasets. For clarity, the sparsity curves report the average results over five runs; the variance follows the same trend as the main comparison and is omitted from the figure to avoid visual clutter. From the experimental results, we can draw the following observations. When the training data proportion decreased from 100% to 20%, the performance of all evaluated models on all metrics significantly declined. This confirms that local structural sparsity will limit the representation learning capability of graph-based recommendation models. Our proposed P2P-IAGR consistently exhibits superior robustness against data sparsity. For instance, on the Food–Kitchen dataset, P2P-IAGR achieves relative improvements of 13.33% in HR@10, and 22.86% in NDCG@10 compared to P2FCDR, respectively. These observations strongly indicate that P2P-IAGR can effectively capture and transfer collaborative signals even when local structural information is highly sparse. By efficiently leveraging cross-domain knowledge to reconstruct the local graph topology, our model successfully alleviates the data sparsity issue and maintains high recommendation accuracy.




Figure 4. Performance comparison of different models with varying proportions of training data on Food-Kitchen and Game-Movie datasets. We vary the training data ratio from 0.2 to 1.0 to evaluate the models' performance under different data sparsity levels. The results illustrate the performance trends of all methods as the amount of available training data decreases.

5.5. Hyper-Parameter Sensitivity (RQ5)

To answer RQ5, we analyze the impact of key hyper-parameters, specifically the number of latent intents K and the contrastive loss weights β and γ .

Impact of K : To investigate the sensitivity of our model to the number of intent prototypes, we vary K from 8 to 128. The results are illustrated in Figure 5. Across all three datasets, we observe a consistent trend, which is that the recommendation performance peaks at $K = 16$. The performance initially improves as K increases, as a larger K allows for a more fine-grained disentanglement of user preferences. However, when K exceeds 32, the performance slightly declines. We attribute this to the fact that an excessively large K may introduce redundant or noisy intents. This redundancy not only causes overfitting but also complicates the cross-domain alignment process. Consequently, we set $K = 16$ as the optimal value for our experiments.




Figure 5. Performance variations with respect to the number of prototypes (K) across three cross-domain scenarios. We vary K in $\{8, 16, 32, 64, 128\}$ to analyze its impact on recommendation accuracy. The curves show a non-linear trend, with the peak performance generally observed around $K = 16$.

To intuitively understand the representation capability of the learned intents, we further visualize the extracted prototypes (using $K = 16$ as an example) in Figure 6. The left panel displays the cosine similarity heatmap among the 16 prototypes. It is evident

that the off-diagonal similarity scores are relatively low (mostly near zero), demonstrating that the learned prototypes are highly orthogonal and successfully capture distinct, non-overlapping semantic preferences. Furthermore, the right panel presents the two-dimensional projection of the principal component analysis of these prototypes. The relatively uniform distribution of the prototypes (marked by red stars) in the latent space confirms that our prototype extraction mechanism effectively avoids pattern collapse.

Impact of β and γ : We investigate the impact of two key hyper-parameters, β (the weight of CCL) and γ (the weight of SCL), on the recommendation performance. The results on *MRR*, *HR@10*, and *NDCG@10* are visualized in Figure 7. The model achieves its best performance across all metrics when $\beta = 0.1$ and $\gamma = 0.1$. This indicates that while contrastive learning signals are crucial for learning better representations, they should act as auxiliary regularizers rather than dominating the primary recommendation objective. Setting the weights too small (e.g., 0.01) leads to insufficient representation learning, whereas setting them too large (e.g., 1.0) causes negative transfer and potential dimensional collapse, which overwhelms the collaborative filtering signals.




Figure 6. Visualization of the extracted intent prototypes ($K = 16$). The left panel shows the cosine similarity heatmap, and the right panel presents the 2D PCA projection. The heatmap displays relatively low off-diagonal similarity values, and the PCA projection illustrates that the prototypes are widely and distinctly distributed across the 2D space.




Figure 7. Parameter sensitivity analysis of β and γ on Food–Kitchen dataset. We vary both β and γ from 0.01 to 1.0. The heatmaps illustrate the joint effect of these two parameters on model performance, with the peak values visibly concentrated in the central region (around 0.1 for both parameters).

5.6. Communication Efficiency and Scalability Analysis

To further assess the scalability of our proposed framework, we additionally conduct experiments on a large-scale cross-domain setting using the Amazon five-core Book and Movie datasets. The Book domain contains 8.89M interactions (603k users, 367k items), and the Movie domain contains 1.69M interactions (123k users, 50k items), with 18,031 overlapping users bridging the two domains. We then compare P2P-IAGR with representative federated baselines (FedGNN and P2FCDR) in terms of transmitted content, communi-

cation cost per round, computational time, and GPU memory overhead. All efficiency experiments were conducted on the same hardware platform, specifically a Linux server equipped with an NVIDIA GeForce RTX 5090 GPU. For fair comparison, the batch size, embedding dimension, and propagation depth were kept consistent with the main experiments. We estimate the communication cost according to the transmitted tensors and empirically measure the training time and GPU memory under this same implementation environment. The reported communication cost includes the transmitted information between two peer clients in one communication round. As shown in Table 4, the general federated graph method suffers from scalability bottlenecks. Specifically, FedGNN requires transmitting full model gradients and computes over the entire graph, leading to large communication costs. For example, each client transmits up to 125.40 MB per round, including uplink and downlink. P2FCDR mitigates this by exchanging overlapping user embeddings, but its communication and memory overhead still scale linearly with the number of shared users, with a complexity of $\mathcal{O}(|\mathcal{U}_{overlap}|D)$. Our P2P-IAGR framework demonstrates superior scalability and communication efficiency. By transmitting only K intent prototypes and their corresponding distributions, the communication cost is bounded to $\mathcal{O}(KD + K^2D)$. Since the number of intents K (e.g., $K = 16$) is smaller than the number of users or parameters (e.g., $K = 16$ with an embedding dimension $D = 128$), P2P-IAGR reduces the communication cost to 0.047 MB per round. This reported value accounts for the bidirectional (uplink and downlink) transmission payload in standard 32-bit float format. Furthermore, because our intent-adaptive graph reconstruction is restricted to sparse neighborhoods via efficient Top- k retrieval, the computational complexity is bounded. Consequently, P2P-IAGR exhibits the lowest training time per epoch (45.6 s) while maintaining low GPU memory consumption (5.8 GB). These experimental results are consistent with our theoretical complexity analysis in Section 3.6 and further demonstrate the scalability of P2P-IAGR on large-scale graphs.

Table 4. Comprehensive efficiency and scalability comparison on the large-scale Amazon Book and Movie cross-domain scenario.

Method	Transmitted Content	Comm. Cost/Round	Time/Epoch	GPU Memory
FedGNN	Model Gradients	125.40 MB	145.2 s	18.5 GB
P2FCDR	Overlapping User Embeds	45.20 MB	89.4 s	12.3 GB
P2P-IAGR (Ours)	K Intent Prototypes	0.047 MB	45.6 s	5.8 GB

Bold indicates the best performance.

6. Conclusions

In this paper, we proposed P2P-IAGR, a novel privacy-preserving peer-to-peer cross-domain collaborative filtering framework, to address the inherent challenges of data sparsity and privacy leakage in recommender systems. To overcome the limitations of coarse-grained preference modeling in existing federated paradigms, we introduced an intent disentanglement mechanism that extracts fine-grained latent prototypes. By securely exchanging these prototypes under Local Differential Privacy (LDP) and aligning them across domains, our approach effectively prevents representation collapse while ensuring privacy protection. Furthermore, we designed an intent-adaptive graph reconstruction module. Guided by the combined cross-domain intent prior, it adaptively augments the sparse local graph and learns robust representations against sparsity. Extensive experiments on real-world datasets demonstrated that P2P-IAGR achieves state-of-the-art performance. In future work, we plan to extend our framework to multi-domain decentralized scenarios and explore more lightweight graph generation methods to further reduce computational overhead on edge devices.

Author Contributions: Conceptualization, M.L.; methodology, M.L.; software, M.L. and H.Z.; validation, M.L., H.Z.; investigation, H.Z.; resources, J.L. and S.C.; data curation, M.L.; writing—original draft preparation, M.L.; writing—review and editing, J.L. and S.C.; visualization, M.L. and H.Z.; supervision, S.C.; project administration, S.C.; funding acquisition, M.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Liaoning Provincial Natural Science Foundation grant number 2024-BS-015, in part by the Fundamental Research Funds for the Central Universities, DMU under Grant 3132026215.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are publicly available. The Amazon dataset can be found at <https://nijianmo.github.io/amazon/index.html> (accessed on 13 July 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Koren, Y.; Bell, R.; Volinsky, C. Matrix factorization techniques for recommender systems. *Computer* **2009**, *42*, 30–37.
2. Rendle, S.; Freudenthaler, C.; Gantner, Z.; Schmidt-Thieme, L. BPR: Bayesian personalized ranking from implicit feedback. *arXiv* **2012**, arXiv:1205.2618.
3. Zhao, C.; Zhao, H.; Li, X.; He, M.; Wang, J.; Fan, J. Cross-domain recommendation via progressive structural alignment. *IEEE Trans. Knowl. Data Eng.* **2023**, *36*, 2401–2415.
4. Zhao, G.; Zhang, X.; Tang, H.; Shen, J.; Qian, X. Domain-oriented knowledge transfer for cross-domain recommendation. *IEEE Trans. Multimed.* **2024**, *26*, 9539–9550.
5. Cao, J.; Lin, X.; Cong, X.; Ya, J.; Liu, T.; Wang, B. Disencdr: Learning disentangled representations for cross-domain recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 267–277.
6. Chen, C.; Wu, H.; Su, J.; Lyu, L.; Zheng, X.; Wang, L. Differential Private Knowledge Transfer for Privacy-Preserving Cross-Domain Recommendation. In *Proceedings of the WWW '22: The ACM Web Conference 2022, Virtual Event, 25–29 April 2022*; Laforest, F., Troncy, R., Simperl, E., Agarwal, D., Gionis, A., Herman, I., Médini, L., Eds.; ACM: New York, NY, USA, 2022; pp. 1455–1465. <https://doi.org/10.1145/3485447.3512192>.
7. Wu, C.; Wu, F.; Cao, Y.; Huang, Y.; Xie, X. FedGNN: Federated Graph Neural Network for Privacy-Preserving Recommendation. *arXiv* **2021**, arXiv:2102.04925.
8. Chen, G.; Zhang, X.; Su, Y.; Lai, Y.; Xiang, J.; Zhang, J.; Zheng, Y. Win-win: A privacy-preserving federated framework for dual-target cross-domain recommendation. *Proc. AAAI Conf. Artif. Intell.* **2023**, *37*, 4149–4156.
9. Lin, Z.; Huang, W.; Zhang, H.; Xu, J.; Liu, W.; Liao, X.; Wang, F.; Wang, S.; Tan, Y. Enhancing Dual-Target Cross-Domain Recommendation with Federated Privacy-Preserving Learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, Republic of Korea, 3–9 August 2024*; pp. 2153–2161. <https://doi.org/10.24963/ijcai.2024/238>.
10. Gao, S.; Wang, S.; Yao, L.; Li, J.; Wang, Z.; Gong, M.; Pan, K. DT-FedSDC: A Dual-Target Federated Framework with Semantic Enhancement and Disentangled Contrastive Learning for Cross-Domain Recommendation. In *CIKM '25: Proceedings of the 34th ACM International Conference on Information and Knowledge Management*; Association for Computing Machinery: New York, NY, USA, 2025; pp. 686–695. <https://doi.org/10.1145/3746252.3761270>.
11. Wang, L.; Wang, S.; Zhang, Q.; Wu, Q.; Xu, M. Federated User Preference Modeling for Privacy-Preserving Cross-Domain Recommendation. *IEEE Trans. Multimed.* **2025**, *27*, 5324–5336. <https://doi.org/10.1109/TMM.2025.3543106>.
12. Chen, S.; Dong, Y.; Shang, Y.; Xu, X.; Lin, X. Federated Privacy-Preserving for Cross-Domain Sequential Recommendation. In *Proceedings of the Artificial Neural Networks and Machine Learning—ICANN 2025*; Senn, W., Sanguineti, M., Saudargiene, A., Tetko, I.V., Villa, A.E.P., Jirsa, V., Bengio, Y., Eds.; Springer: Cham, Switzerland, 2026; pp. 41–52.
13. Zhu, F.; Wang, Y.; Chen, C.; Zhou, J.; Li, L.; Liu, G. Cross-Domain Recommendation: Challenges, Progress, and Prospects. *arXiv* **2021**, arXiv:2103.01696.
14. Cao, J.; Cong, X.; Sheng, J.; Liu, T.; Wang, B. Contrastive cross-domain sequential recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 138–147.

15. Wang, Y.; Feng, C.; Guo, C.; Chu, Y.; Hwang, J. Solving the Sparsity Problem in Recommendations via Cross-Domain Item Embedding Based on Co-Clustering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, Australia, 11–15 February 2019*; Culpepper, J.S., Moffat, A., Bennett, P.N., Lerman, K., Eds.; ACM: New York, NY, USA, 2019; pp. 717–725. <https://doi.org/10.1145/3289600.3290973>.
16. Zhao, C.; Li, C.; Xiao, R.; Deng, H.; Sun, A. CATN: Cross-Domain Recommendation for Cold-Start Users via Aspect Transfer Network. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2020, Virtual Event, 25–30 July 2020*; Huang, J.X., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., Liu, Y., Eds.; ACM: New York, NY, USA, 2020; pp. 229–238. <https://doi.org/10.1145/3397271.3401169>.
17. Lian, J.; Zhang, F.; Xie, X.; Sun, G. CCCFNet: A Content-Boosted Collaborative Filtering Neural Network for Cross Domain Recommender Systems. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, 3–7 April 2017*; Barrett, R., Cummings, R., Agichtein, E., Gabrilovich, E., Eds.; ACM: New York, NY, USA, 2017; pp. 817–818. <https://doi.org/10.1145/3041021.3054207>.
18. He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; Chua, T.S. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*; ACM: New York, NY, USA, 2017; pp. 173–182.
19. Zhu, F.; Wang, Y.; Chen, C.; Liu, G.; Orgun, M.A.; Wu, J. A Deep Framework for Cross-Domain and Cross-System Recommendations. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, Stockholm, Sweden, 13–19 July 2018*; Lang, J., Ed.; International Joint Conferences on Artificial Intelligence Organization: California, CA, USA, 2018; pp. 3711–3717. <https://doi.org/10.24963/IJCAI.2018/516>.
20. Hu, G.; Zhang, Y.; Yang, Q. CoNet: Collaborative Cross Networks for Cross-Domain Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, 22–26 October 2018*; Cuzzocrea, A., Allan, J., Paton, N.W., Srivastava, D., Agrawal, R., Broder, A.Z., Zaki, M.J., Candan, K.S., Labrinidis, A., Schuster, A., et al., Eds.; ACM: New York, NY, USA, 2018; pp. 667–676. <https://doi.org/10.1145/3269206.3271684>.
21. Zhu, F.; Wang, Y.; Chen, C.; Liu, G.; Zheng, X. A Graphical and Attentional Framework for Dual-Target Cross-Domain Recommendation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*; Bessiere, C., Ed.; International Joint Conferences on Artificial Intelligence Organization: California, CA, USA, 2020; pp. 3001–3008. <https://doi.org/10.24963/IJCAI.2020/415>.
22. Li, P.; Tuzhilin, A. DDTCDR: Deep Dual Transfer Cross Domain Recommendation. In *Proceedings of the WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, 3–7 February 2020*; Caverlee, J., Hu, X.B., Lalmas, M., Wang, W., Eds.; ACM: New York, NY, USA, 2020; pp. 331–339. <https://doi.org/10.1145/3336191.3371793>.
23. Zhu, F.; Wang, Y.; Chen, C.; Liu, G.; Orgun, M.; Wu, J. A deep framework for cross-domain and cross-system recommendations. *arXiv* **2020**, arXiv:2009.06215.
24. Chen, X.; Zhang, Y.; Tsang, I.W.; Pan, Y.; Su, J. Towards Equivalent Transformation of User Preferences in Cross Domain Recommendation. *arXiv* **2020**, arXiv:2009.06884.
25. Zhao, C.; Li, C.; Fu, C. Cross-Domain Recommendation via Preference Propagation GraphNet. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, 3–7 November 2019*; Zhu, W., Tao, D., Cheng, X., Cui, P., Rundensteiner, E.A., Carmel, D., He, Q., Yu, J.X., Eds.; ACM: New York, NY, USA, 2019; pp. 2165–2168. <https://doi.org/10.1145/3357384.3358166>.
26. Cao, J.; Sheng, J.; Cong, X.; Liu, T.; Wang, B. Cross-Domain Recommendation to Cold-Start Users via Variational Information Bottleneck. In *Proceedings of the 38th IEEE International Conference on Data Engineering, ICDE 2022, Kuala Lumpur, Malaysia, 9–12 May 2022*; IEEE: Piscataway, NJ, USA, 2022; pp. 2209–2223. <https://doi.org/10.1109/ICDE53745.2022.00211>.
27. Pei, J.; Liu, W.; Li, J.; Wang, L.; Liu, C. A review of federated learning methods in heterogeneous scenarios. *IEEE Trans. Consum. Electron.* **2024**, *70*, 5983–5999.
28. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Y Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017 PMLR, Fort Lauderdale, FL, USA, 20–22 April 2017*; Singh, A., Zhu, X.J., Eds.; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2017; pp. 1273–1282.
29. Chai, D.; Wang, L.; Chen, K.; Yang, Q. Secure Federated Matrix Factorization. *IEEE Intell. Syst.* **2021**, *36*, 11–20. <https://doi.org/10.1109/MIS.2020.3014880>.
30. Feng, C.; Feng, D.; Huang, G.; Liu, Z.; Wang, Z.; Xia, X.G. Robust privacy-preserving recommendation systems driven by multimodal federated learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *36*, 8896–8910.
31. Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; Zhang, C. FedProto: Federated Prototype Learning across Heterogeneous Clients. *Proc. Aaa Conf. Artif. Intell.* **2022**, *36*, 8432–8440. <https://doi.org/10.1609/AAAI.V36I8.20819>.

32. He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; Wang, M. 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*; Association for Computing Machinery: New York, NY, USA, 2020; pp. 639–648.
33. Ye, Z.; Zhang, X.; Chen, X.; Xiong, H.; Yu, D. Adaptive Clustering Based Personalized Federated Learning Framework for Next POI Recommendation With Location Noise. *IEEE Trans. Knowl. Data Eng.* **2024**, *36*, 1843–1856. <https://doi.org/10.1109/TKDE.2023.3312511>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.