Win-Win: A Privacy-Preserving Federated Framework for Dual-Target Cross-Domain Recommendation

Gaode Chen^{1, 2, 3, 4}*, Xinghua Zhang^{1, 2}, Yijun Su^{3, 4†}, Yantong Lai^{1, 2}, Ji Xiang^{1, 2}, Junbo Zhang^{3, 4†}, Yu Zheng^{3, 4}

 ¹ Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
² School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China
³ JD iCity, JD Technology, Beijing, China
⁴ JD Intelligent Cities Research, Beijing, China
{chengaode, zhangxinghua, laiyantong, xiangji}@iie.ac.cn, suyijun.ucas@gmail.com, {msjunbozhang, msyuzheng}@outlook.com

Abstract

Cross-domain recommendation (CDR) aims to alleviate the data sparsity by transferring knowledge from an informative source domain to the target domain, which inevitably proposes the stern challenges to data privacy and transferability during the transfer process. A small amount of recent CDR works have investigated the privacy protection, while they still suffer from satisfying practical requirements (e.g., limited privacy-preserving ability) and preventing the potential risk of negative transfer. To address above challenging problems, we propose a novel and unified privacy-preserving federated framework for dual-target CDR, namely P2FCDR. We design P2FCDR as peer-to-peer federated network architecture to ensure the local data storage and privacy protection of business partners. Specifically, for the special knowledge transfer process in CDR under federated settings, we initialize an optimizable orthogonal mapping matrix to learn the embedding transformation across domains and adopt the local differential privacy technique on the transformed embedding before exchanging across domains, which provides the more reliable privacy protection. Furthermore, we exploit the similarity between in-domain and cross-domain embedding, and develop a gated selecting vector to refine the information fusion for more accurate dual-transfer. Extensive experiments on three real-world datasets demonstrate that P2FCDR significantly outperforms the state-of-the-art methods and effectively protects data privacy.

Introduction

In the era of information explosion, recommender systems are instrumental to alleviate information overload. Especially, Collaborative filtering (Wang et al. 2019a) (CF) is a promising and widely used technique for modern recommender systems, which learns the latent embedding of users and items effectively, and then performs the recommendation based on these embedding vectors. However, these methods inevitably suffer from the long-standing data sparsity problem to some extent. This problem makes it hard to model user preference accurately and efficiently, which can lead to a dramatic decrease in recommendation quality.

Cross Domain Recommendation (CDR) has been popularly studied recently to alleviate the data sparsity problem, which leverages the data in an auxiliary domain to improve the recommendation performance of the target domain. Although existing CDR methods (Liu et al. 2020a; Li and Tuzhilin 2020) have yielded immense success, they usually assume that user behavior data can be fully shared across domains, which undoubtedly poses a significant risk of privacy leakage since user data are very sensitive. Several works have focused on studying privacy protection for CDR, but they still suffer from satisfying practical requirements, which are mainly embodied in three aspects:

- Limited Privacy-preserving Ability. Existing privacypreserving CDR approaches (Gao et al. 2019b,a, 2021; Chen et al. 2022) share either item embedding or differential private interaction matrix with target domain to protect privacy. The plaintext embedding is still at risk of being inferred, and sharing original interaction data has become an unacceptable behavior as new privacy regulations (e.g., GDPR) are being enacted around the world. So existing privacy-preserving CDR methods are limited by current privacy-preserving criteria.
- **Incomplete Scenario Coverage.** The existing privacypreserving CDR models (Liu et al. 2021; Yan et al. 2022) are mainly proposed to protect the data privacy of individual customers. These methods treat edge devices of users as clients, and each client locally stores the private data of a single user. As is known to all, a company tends to seek cooperation with another one in a different field, and then enhances business capability through these complementary information. Therefore, the privacy protection of business partners is vital but existing methods are short of exploration for this scenario.
- Unidirectional Benefit Target. The privacy-preserving CDR methods (Chen et al. 2022; Yan et al. 2022) men-

^{*}The paper was done when Gaode Chen was an intern at JD Intelligent Cities Research under the supervision of the Junbo Zhang. *Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tioned above focus on the unidirectional transfer for improving the target domain performance with the aid of source domain information. Thus, the company furnishing the source domain data benefits less from the unidirectional transfer, and is hard to be motivated to participate this process. Furthermore, existing privacypreserving CDR methods cannot work well by simply changing the transfer direction from the target domain to the source domain, which may sustain negative transfer and the risk of potential privacy leakage. Thus, devising a unified privacy-preserving dual-target CDR framework for bidirectional transfer can improve the enthusiasm of all participants and is more in line with real application.

In addition to being restricted by above practical requirements, another challenge in CDR is how to extract useful signal from one domain to improve another. As not all features contained in the source domain can be conducive to the target. Most of the existing CDR methods (Liu et al. 2020b; Li and Tuzhilin 2020) conduct feature alignment and fusion based on the entire in-domain and cross-domain embedding representations, which may contribute to negative transfer due to coarse information combination. Thus, extracting fine-grained signal from the transferred information and digging further into the process of feature fusion plays a pivotal role in improving the performance of CDR.

To tackle above challenges, we design a novel and unified Privacy-Preserving Federated framework for dual-target Cross-Domain Recommendation, called **P2FCDR**. We aim to develop a more practical solution for CDR, and our framework has advantages in following three aspects. Firstly, we adopt the federated learning mechanism (Yang et al. 2019) to ensure that the original data are kept locally and synchronize the updated model parameters in domains. Besides, our framework is designed as a peer-to-peer federated network architecture to further meet the privacy-preserving criteria, which ensures that there are no curious or malicious third parties. Since the CDR scenario has a necessary knowledge transfer process, we further utilize local differential privacy (Choi et al. 2018) (LDP) to perturb user embedding before exchanging across domains to prevent external attackers and another participant from inferring privacy-sensitive information, which is different from prior works that share plaintext item embedding or original data with the target domain. Secondly, P2FCDR takes business partners as clients and satisfies the scenario of cooperation between companies. Thirdly, our framework is symmetric and improves the recommendation performance simultaneously in both domains.

Specifically, we first use an embedding-based recommendation model in each domain to obtain user and item embedding. Inspired by (Li and Tuzhilin 2020), we implement the transformation of user embedding between domains using an orthogonal mapping matrix and employ LDP to hide the real distribution of the transformed embedding. Then, we introduce a gated selecting vector in each domain to refine the fusion process between in-domain user embedding and cross-domain user embedding (transformed embedding from another domain) at feature level, which can extract useful signal that is highly related to the target domain, further avoiding negative transfer. Finally, we synchronize the updated model parameters between the two domains.

The contributions of our work are summarized as follows:

- To the best of our knowledge, **P2FCDR** is the first privacy-preserving federated framework for dual-target cross-domain recommendation, which is more suitable for practical needs from the perspectives of privacy protection, scenario coverage, and bidirectional motivating.
- We explicitly consider information fusion in CDR and derive a gated selecting vector, aiming at extracting the fine-grained signal at feature level that is highly related to the target domain to avoid negative transfer.
- We conduct extensive experiments on three real-world datasets to confirm the effectiveness of our proposed method in terms of recommendation performance and privacy protection.

Related Work

Cross-Domain Recommendation

Cross-Domain Recommendation (CDR) technique is an effective way of alleviating the data sparsity issue in recommender systems by leveraging the knowledge from relevant domains, which is well summarized and classified in a survey (Zhu et al. 2021). Existing studies can be classified into three categories according to the relevance between domains: user-level relevance, item-level relevance, and content-level relevance. Meanwhile, the scenarios of CDR can be identified as single-target CDR (Wang et al. 2019b; Zhao et al. 2020), multi-domain recommendation (Zhang et al. 2016), dual-target CDR (Krishnan et al. 2020).

Our work falls into category of user-level relevance and category of dual-target CDR. Specifically, user-level relevance in (Zhu et al. 2021) is defined as the multiple domains that have common users and different levels of items. Since most existing CDR models assume that data across domains are accessible directly, ignoring the privacy issues, we aim to solve the privacy issues in CDR based on federated learning and local differential privacy.

Privacy-Preserving Recommendation

As one of the most popular personalized services in today's online systems, recommendation systems are close to users' privacy data such as check-in and purchase data. However, it is more necessary to avoid privacy leakage in the CDR scenario, because it has a process of knowledge transfer. Recently, the improvement of relevant laws and regulations has also pushed the development of privacypreserving recommendation. Both CCMF (Gao et al. 2019b) and PriCDR (Chen et al. 2022) add noise to the original user-item matrix based on differential privacy to preserve privacy and directly share it with the target domain. Besides, NATR (Gao et al. 2019a, 2021) shares the item embedding with the target domain in plaintext to avoid privacy leakage, while there is still at risk of being attacked. Recently, FedCT (Liu et al. 2021) and FedCDR (Yan et al. 2022) introduce federated learning into CDR to better protect user



Figure 1: The overall framework of the proposed P2FCDR, which consists of Representation Modeling, Embedding Transformation, Embedding Fusion, and Parameter Update and Synchronization. In addition, there are two cross-domain communications in each iteration, as shown by the red and dark arrows.

data privacy. They both perform CDR on the user's edge device and only improve the recommendation performance of the target domain.

In summary, the existing privacy-preserving recommendation methods still have limitations in satisfying practical situations: *Limited Privacy-preserving Ability, Incomplete Scenario Coverage*, and *Unidirectional Benefit Target*. In this work, we implement a privacy-preserving CDR framework in more practical settings.

Methodology

Problem Definition

We consider two domains \mathcal{A} and \mathcal{B} , which have the same set of users U (of size m = |U|) but different user-item pairs. Let the set of items in \mathcal{A} and \mathcal{B} be V_A (of size $n_a = |V_A|$) and V_B (of size $n_b = |V_B|$), respectively. Let $\mathbf{R}^A \in \mathbb{R}^{m \times n_a}$ $(\mathbf{R}^B \in \mathbb{R}^{m \times n_b}, \text{ resp.})$ represents the user-item interaction matrix of \mathcal{A} (\mathcal{B} , resp.) from explicit feedback, such as ratings, or implicit feedback, such as clicks. We use \mathbf{u}_i^A and \mathbf{u}_i^B to denote the latent embedding of the user *i* in \mathcal{A} and \mathcal{B} , respectively. We consider Top-N recommendation in each domain that learning a function to estimate the scores of unobserved entries in the interaction matrix, which are later used for ranking. The goal of CDR is to improve the recommendation performance in both domains while not sharing original interaction data. We do not distinguish a source domain or a target domain since the recommendation task for each domain is performed in an unified framework here.

Overview

We propose P2FCDR, a privacy-preserving federated framework for the dual-target cross-domain recommendation. Our framework is designed as peer-to-peer federated network architecture, where user-item interaction data are stored locally without third-party involvement to further protect data privacy. Furthermore, we use local differential privacy for the transferred embedding to avoid violating the business privacy policy. Figure 1 illustrates our P2FCDR framework, which consists of the following four modules and involves two communications across domains in one iteration.

- **Representation Modeling.** We obtain the user and item representations in each domain with data stored locally.
- Embedding Transformation. We transfer embedding of common users in domain A into domain B and protect the transferred embedding with local differential privacy, and vice versa for domain B.
- Embedding Fusion. We employ a gated selecting vector in each domain to refine the information fusion of indomain and cross-domain embedding at feature level.
- **Parameter Update and Synchronization.** We update the parameters based on the recommendation loss in each domain and synchronize corresponding parameters.

In the following, we elaborate in detail on these four modules and analyze privacy for the entire process.

Representation Modeling

For the user interaction data of each domain, through deep neural network based recommendation models, e.g., Deep Matrix Factorization (DMF) (Xue et al. 2017) and Neural Matrix Factorization (NeuMF) (He et al. 2017), we generate the embedding representations for users and items in each domain. In this paper, we take DMF as an example and show representation modeling in domain A.

In domain \mathcal{A} , each user *i* is represented as the *i*-th row of matrix $\mathbf{R}^{\mathcal{A}}$, i.e. $\mathbf{R}_{i*}^{\mathcal{A}}$. Similarly, each item *j* is represented as

the *j*-th column of matrix \mathbf{R}^A , i.e. \mathbf{R}_{*j}^A . We adopt two multilayer networks to map user *i* and item *j* to a low-dimensional embedding vector in a latent space, respectively.

$$\mathbf{u}_{i}^{A} = \sigma \left(\cdots \sigma \left(\mathbf{W}_{U_{2}}^{A} \sigma \left(\mathbf{W}_{U_{1}}^{A} \mathbf{R}_{i*}^{A} \right) \right) \right) \in \mathbb{R}^{k}$$

$$\mathbf{v}_{j}^{A} = \sigma \left(\cdots \sigma \left(\mathbf{W}_{V_{2}}^{A} \sigma \left(\mathbf{W}_{V_{1}}^{A} \mathbf{R}_{*j}^{A} \right) \right) \right) \in \mathbb{R}^{k}$$
(1)

where σ is the non-linear activation ReLU, k is the dimension of user and item embedding, \mathbf{W}_{U_1} , \mathbf{W}_{U_2} , ... and \mathbf{W}_{V_1} , \mathbf{W}_{V_2} , ... are the weights of different layers in multi-layer networks of users and items, respectively. Analogously, we can derive \mathbf{u}_i^B and \mathbf{v}_z^B for domain \mathcal{B} , where item $z \in V_B$.

Embedding Transformation

To obtain better understanding of user preference in domain \mathcal{A} , we utilize external user embedding from domain \mathcal{B} and combine them together, and vice versa. However, the data distribution may vary considerably across domains, we need to learn the mapping relationship between the feature spaces of the two domains to achieve better domain adaptation. Thus, the embedding of the common users in the auxiliary domain is first transformed and then sent to the target domain, which involves the first cross-domain communication.

Inspired by (Li and Tuzhilin 2020), we introduce a latent orthogonal matrix for transforming user embedding from the auxiliary domain to the target domain, e.g., from domain \mathcal{A} to domain \mathcal{B} . Specifically, the orthogonal mapping matrix has the following two advantages in our scenario. First, it preserves similarities between user embedding before and after transformation, since the orthogonal transformation preserves the inner product of vectors. Second, its inverse mapping matrix is equivalent to its transpose. Thus, the transformation of domain \mathcal{B} to domain \mathcal{A} can directly use its transpose to simplify the learning procedure and reduce the computation complexity.

We adopt local differential privacy (LDP) to further hide the transformed user representation before transferring to another domain. It not only prevents external attackers from inferring attacks (Chai et al. 2020), but also prevents another domain (business partner) from inferring the real user representations based on the orthogonal mapping matrix, because this may violate commercial privacy. Thus, the embedding \mathbf{u}_i^A of user *i* in domain \mathcal{A} transferred to domain \mathcal{B} can be represented as $\mathbf{u}_i^{AB} \in \mathbb{R}^k$, which can be defined as follows. Similarly, we can also obtain $\mathbf{u}_i^{BA} \in \mathbb{R}^k$.

$$\mathbf{u}_{i}^{AB} = \mathbf{u}_{i}^{A} \mathbf{X}^{A} + La\left(0,\lambda\right)$$
$$\mathbf{u}_{i}^{BA} = \mathbf{u}_{i}^{B} \left[\mathbf{X}^{B}\right]^{\top} + La\left(0,\lambda\right)$$
(2)

where $\mathbf{X}^A, \mathbf{X}^B \in \mathbb{R}^{k \times k}$ are orthogonal mapping matrices in domain \mathcal{A} and \mathcal{B} , respectively, and are consistently synchronized and equal. $[\cdot]^{\top}$ represents the transpose of the matrix. $La(0, \lambda)$ denotes Laplace noise with a mean value of 0, and λ controls the strength of the Laplace noise. The larger λ , the greater the noise and the better for privacy protection.

Embedding Fusion

In prior studies, the target domain in CDR usually directly combines the transferred user embedding from the auxiliary domain, e.g., Concatenation, Max-Pooling and Average-Pooling. If the data sparsity problem also exists in the auxiliary domain, the transferred embedding is not sufficient, or the embedding in the target domain is already sufficient and does not require extra supplement, the negative transfer will occur. So feature-level selecting vector is in demand to refine the information fusion process of the transferred embedding for extracting fine-grained signal that is highly related to the target domain. We take domain \mathcal{A} as an example and use the L_1 distance to represent the similarity between \mathbf{u}_i^A and \mathbf{u}_i^{BA} , the same is true for domain \mathcal{B} :

$$\mathbf{\Delta}_{i}^{A} = \left| \mathbf{u}_{i}^{A} - \mathbf{u}_{i}^{BA} \right| \tag{3}$$

Ideally, a specific dimension of Δ_i^A denotes the similarity between \mathbf{u}_i^A and \mathbf{u}_i^{BA} at the feature level, and the smaller the value is, the closer it is to the target. Meanwhile, the distances between the in-domain user embedding and corresponding cross-domain embedding in each feature dimension contribute differently to the target. Thus, a two-layer fully-connected neural network is utilized to automatically learn and derive the gated selecting vector $\mathbf{s}_i^A \in \mathbb{R}^k$ of Δ_i^A in domain \mathcal{A} for controlling information flow in each dimension as follows, so was domain \mathcal{B} :

$$\mathbf{s}_{i}^{A} = \text{Sigmoid} \left(\mathbf{W}_{s_{2}}^{A} \left(\sigma \left(\mathbf{W}_{s_{1}}^{A} \mathbf{\Delta}_{i}^{A} + \mathbf{b}_{s_{1}}^{A} \right) + \mathbf{b}_{s_{2}}^{A} \right) \right) \quad (4)$$

where the $\mathbf{W}_{s_1}^A$, $\mathbf{W}_{s_2}^A$, $\mathbf{b}_{s_1}^A$ and $\mathbf{b}_{s_2}^A$ are trainable weight and bias, and σ denotes the non-linear activation function ReLU.

The refined representation $\tilde{\mathbf{u}}_i^A$ of user *i* in domain \mathcal{A} should be obtained through the weighted combinations of the in-domain representations \mathbf{u}_i^A and the cross-domain representations \mathbf{u}_i^{BA} by the above gated selecting vector \mathbf{s}_i^A :

$$\tilde{\mathbf{u}}_{i}^{A} = \mathbf{s}_{i}^{A} \cdot \mathbf{u}_{i}^{A} + \left(\mathbf{1} - \mathbf{s}_{i}^{A}\right) \cdot \mathbf{u}_{i}^{BA}$$
(5)

Parameter Update and Synchronization

After the aforementioned operation, our problem further turns to predict user interaction in each domain based on the refined user embedding vector $\tilde{\mathbf{u}}_i^A$ and the item embedding vector \mathbf{v}_j^A (Taking domain \mathcal{A} as an example, so was domain \mathcal{B}). Here we adopt a simple but widely-used inner product model, to estimate the value of \hat{y}_{ij}^A , which is the interaction probability between a given pair of user and item.

$$\hat{y}_{ij}^{A} = \text{Sigmoid}\left(\left[\tilde{\mathbf{u}}_{i}^{A}\right]^{\top} \mathbf{v}_{j}^{A}\right) \tag{6}$$

Following the previous works (Xue et al. 2017; Xie et al. 2021), we train our model in both domains with the following objective function:

$$\mathcal{L}(y_{ij}, \hat{y}_{ij}) = -\sum_{(i,j)\in Y^+\cup Y^-} \left(\frac{y_{ij}}{\max(R)}\log\hat{y}_{ij} + \left(1 - \frac{y_{ij}}{\max(R)}\right)\log(1 - \hat{y}_{ij})\right)$$
(7)

where Y^+ denotes the set of observed interactions and Y^- denotes negative instances sampled from the unobserved interactions. We use the max(R) for normalization which

is the max score in a dataset. For explicit feedback, e.g. $\max(R)$ is 5 in a 5-star system, different values of y_{ii} have different influences to the loss. For implicit feedback, $\max(R)$ is 1, and y_{ij} is 0 or 1.

Using only one mapping matrix is not suitable for practical deployment scenarios. In a real situation, the models of domain \mathcal{A} and \mathcal{B} are respectively deployed in different and independent places, they cannot optimize the same orthogonal mapping matrix simultaneously. Therefore, after updating the model parameters in each domain according to the above loss function \mathcal{L} , we need to synchronize the orthogonal mapping matrices in the two domains for the next iteration, which involves the second cross-domain communication. Taking domain A as an example, domain B directly sends its updated orthogonal mapping matrix $\mathbf{X}^{B\prime}$ to domain \mathcal{A} in plaintext, and domain \mathcal{A} conducts Average Pooling on the $\mathbf{X}^{B\prime}$ and in-domain updated orthogonal mapping matrix $\mathbf{X}^{A\prime}$ to obtain a new orthogonal mapping matrix \mathbf{X}^{A} . The same is true for domain \mathcal{B} , which ensures that the orthogonal mapping matrix in the two domains are consistently the same in each iteration.

Privacy Analysis

In this section, we mainly focus on the privacy analysis of our framework P2FCDR. Our work aims to avoid the leakage of user privacy information, which is an important concern in existing CDR approaches.

In our proposed framework, user data on a business platform never leave local but are always stored on their devices, e.g., databases. We design the federated learning architecture as a peer-to-peer network structure to further reduce the risk of privacy leakage, i.e., there are no curious or malicious third parties. The two domains communicate only twice in one iteration to exchange model-related information. We analyze each communication one by one as follows.

In the first communication, the transferred user embedding is exchanged between the two domains, but doing so directly has the potential to infer the real data of users (Chai et al. 2020). Thus, we employ the LDP technique to hide user representations by adding noise that obeys the Laplace distribution before transferring transformed user embedding to another domain, which further enhances the privacy protection of user data. Because this not only prevents external attackers from intercepting the transferred information, but also another business partner cannot restore the original user embedding through the orthogonal mapping matrix, which violates the business privacy policy. In LDP, a higher noise strength λ leads to a decrease in data availability, which will reduce the recommendation performance. Therefore, we need to strike a trade-off between recommendation performance and privacy protection.

In the second communication, the updated orthogonal mapping matrices in each domain are exchanged with each other. We do not use privacy-preserving methods for this communication. Because the orthogonal mapping matrix represents the transfer correspondence between two domains, which not involve sensitive data. Furthermore, even if an external attacker intercepts the updated orthogonal map-

Datasets	Domain	Users	Items	Ratings	Sparsity
Movie & Book	Movie Book	14,591	11,270 19,977	397,034 428,493	99.76% 99.85%
Music & Movie	Music Movie	8,252	8,737 9,218	196,792 284,589	99.73% 99.63%
Book & Music	Book Music	2,424	3,251 2,880	57,404 53,498	99.27% 99.23%

Table 1: The statistics of datasets.

ping matrix and the transferred user embedding after differential privacy, it is hard to infer valid information.

Experiments

Experimental Setup

Datasets We study the effectiveness of our P2FCDR on three largest domains on a real-world public dataset Amazon¹, i.e., Movies and TV (Movie), Books (Book), and CD Vinyl (Music). Following (Chen et al. 2020, 2022), we make a pairwise combinations amongst the three domains and only choose the user-item interactions of the common users across domains, i.e., Movie & Book, Music & Movie, and Book & Music. For the data in these three couple datasets, we first transform them into implicit data, where each entry is marked as 0 or 1, indicating whether the user has rated the item. Then, we filter the datasets to retain users with number of ratings greater than 5 and items with number of ratings greater than 10. Table 1 summarizes the detailed statistics of the three couple datasets.

Comparison methods We compare P2FCDR with the following three categories of methods: Single-Domain Recommendation, Cross-Domain Recommendation, and Privacy-Preserving Cross-Domain Recommendation. **Single-Domain Recommendation:**

- NeuMF (He et al. 2017) is a neural network based model that replaces the conventional inner product with a neural architecture to improve recommendation accuracy.
- DMF (Xue et al. 2017) is a deep neural network based recommendation model that employs a deep architecture to learn the low-dimensional factors of users and items.

Cross-Domain Recommendation:

- CoNet (Hu, Zhang, and Yang 2018) is a collaborative cross network that enables knowledge transfer across domains by cross connections between base networks.
- DDTCDR (Li and Tuzhilin 2020) is a deep dual transfer network that transfers knowledge with orthogonal transformation across domains.
- BiTGCF (Liu et al. 2020b) is a CDR model based on a graph collaborative filtering network to fuse users' common features and domain-specific features.
- ETL (Chen et al. 2020) is a recent state-of-the-art CDR model that adopts a equivalent transformation to model the joint distribution of user behaviors across domains.

¹http://jmcauley.ucsd.edu/data/amazon/

	$Movie \leftrightarrow Book$			$\mathbf{Music}\leftrightarrow\mathbf{Movie}$			$Book \leftrightarrow Music$					
	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10
Single-Domain Recommendation												
NeuMF	0.5094	0.3008	0.4746	0.2568	0.4476	0.2366	0.4295	0.2395	0.3787	0.1959	0.4191	0.2318
DMF	0.4919	0.2823	0.4938	0.2874	0.4584	0.2534	0.4439	0.2499	0.3651	0.1943	0.4295	0.2410
Cross-Domain Recommendation												
CoNet	0.3501	0.1884	0.2878	0.1685	0.2248	0.1166	0.2793	0.1471	0.3148	0.1353	0.3548	0.1893
DDTCDR	0.5436	0.3248	0.5767	0.3390	0.5148	0.3211	0.4838	0.2721	0.4201	0.2369	0.4548	0.2475
BiTGCF	0.5715	0.3448	0.6377	0.4042	0.6195	0.3764	0.5383	0.3201	0.5379	0.3108	0.5255	0.3057
ETL	0.5756	0.3591	0.6362	0.3906	0.6119	0.3758	0.5326	0.3205	0.4319	0.2465	0.4600	0.2475
Privacy-Preserving Cross-Domain Recommendation												
PriCDR-S		-	0.6029	0.3504	-	-	0.5319	0.2982	-	-	0.5321	0.3089
FedCDR	-	_	0.6094	0.3756	-	-	0.5247	0.3101	-	_	0.5057	0.2902
P2FCDR	0.6123	0.3754	0.6793	0.4299	0.6528	0.4059	0.5915	0.3503	0.5689	0.3317	0.5470	0.3231

Table 2: Comparisons with baselines on three pairs of datasets. The best performance in each column is bolded number.

Privacy-Preserving Cross-Domain Recommendation:

- **PriCDR-S** (Chen et al. 2022) is a privacy-preserving CDR model designed for business partners scenario, which shares the differentially private rating matrix from source domain with the target domain.
- **FedCDR** (Yan et al. 2022) is a privacy-preserving federated CDR model designed for individual customers scenario, which learns the cross-domain embedding transformation model on the server side.

Note that both PriCDR-S and FedCDR are single-target CDR frameworks, so we can only show their recommendation performance on one of the domains.

Evaluation method To evaluate the recommendation performance, we use the leave-one-out method which is widely used in the literature of recommendation (He et al. 2017). Specifically, we held out the latest interaction as the test set and utilized the remaining data for training. Then, we follow the common strategy which randomly samples 99 (negative) items that are not interacted with by the user and then evaluate how well the recommender can rank the test item against these negative ones. Since we aim at Top-*N* item recommendation, the typical evaluation metrics are hit ratio (HR), normalized discounted cumulative gain (NDCG), where the ranked list is cut off at 10 in our experiments. HR measures whether the test item is ranked on the Top-*N* list while NDCG measures the specific ranking quality that assigns high scores to the hits at top position ranks.

Parameter settings For the representation modeling of users and items, we both use a two-layer fully connected network with dimensions 64 and 64 respectively, and obtain the final embedding dimension k as 64. Considering the trade-off between recommendation performance and privacy protection, we set λ to 0.02. For the learning of gated selecting vector, we use a two-layer fully connected network with dimension 64 and 64, respectively. When training our models, we choose Adam as the optimizer, and set the learning rate to 0.001. Meanwhile, we select a batch of users according to the IDs of the common user to construct mini-batches, and set the batch size to 256.

Performance Evaluation

In this section, we report the recommendation performance of different methods and discuss the findings. Table 2 shows the summarized results of our experiments on the three couple datasets in terms of two metrics, HR@10 and NDCG@10. We have the following observations:

- For those single-domain methods that are only trained with in-domain interaction data, they suffer from the data sparsity problem, thus resulting in relatively poor recommendation performance, especially on Book & Music with sparse data and smaller data volumes. Our proposed method P2FCDR consistently outperforms single domain baselines in both evaluation metrics, especially achieving higher recommendation performance on sparser domains, such as the Book domain of Movie & Book. This indicates the proposed P2FCDR is a promising solution for the CDR task, and can effectively alleviate the data sparsity problem.
- Cross-domain recommendation methods generally outperform the single-domain methods, indicating the effectiveness of transferring knowledge across domains in the recommendation for data sparsity issue. In particular, BiTGCF takes into account domain-specific features and ETL also models both the overlapped and domain-specific features, thus they show better performance. While DDTCDR does not have optimization for refining feature fusion, so the performance is relatively poor. The performance of Conet is not satisfactory, and there are similar results in (Chen et al. 2020), because its learning mechanism breaks the joint behavior pattern in CDR. P2FCDR performs even better than those crossdomain methods with the risk of leaking user privacy. This demonstrates that P2FCDR can effectively extract signal highly relevant to the target domain in transferred embedding under the premise of privacy protection.
- Since not yet privacy-preserving CDR methods designed for dual-target, we can only demonstrate their recommendation performance on a single domain (as target domain). PriCDR-S mainly focuses on the embedding alignment between user embedding vectors in the auxiliary and target domains, FedCDR also focuses on learning the mapping relationship between user embedding

	Movie \leftrightarrow Book						
	HR@10	NDCG@10	HR@10	NDCG@10			
w/o ET	0.5710	0.3557	0.6493	0.4053			
w/o EF	0.5688	0.3387	0.6158	0.3760			
P2FCDR	0.6123	0.3754	0.6793	0.4299			

Table 3: Ablation studies of our method on Movie & Book dataset.

vectors in the two domains and is more suitable for 2C scenarios, making them fail to beat the most competitive CDR method. Meanwhile, P2FCDR utilizes federated learning and LDP to effectively protect data privacy and works well in transferring knowledge from auxiliary domain, showing better recommendation performance than privacy-preserving CDR methods.

Ablation Study

We further investigate that both orthogonal matrix based embedding transformation and feature level based embedding fusion are essential parts of our framework P2FCDR. We present the results of ablation experiments on Movie & Book in Table 3, we can observe that: (1) Without Embedding Transformation (w/o ET), we directly transfer the differentially private user representations from the auxiliary domain to the target domain. Embedding transformation contributes to 6.75% and 4.42% relatively increase for HR@10 (5.25% and 5.72% for NDCG@10) on Movie and Book domain, respectively, which indicates the orthogonal matrix can effectively learn the transfer correspondence across domains and improve recommendation performance. (2) Without Embedding Fusion (w/o EF), we utilize Average Pooling for the combination of in-domain and cross-domain user embedding instead of based on the gated selecting vector. The refining of user embedding at feature level improves the relative 7.10% and 9.35% recommendation performance for HR@10 (9.78% and 12.54% for NDCG@10) on Movie and Book, respectively. The performance improvement it brings is more obvious in sparser domain, e.g., Book domain of Movie & Book. This confirms the utility of gated selecting vector in distilling signal from the transferred user embedding and enhancing user representations in target domain.

Security vs Performance

To further explore the impact of privacy-preserving technologies, i.e., LDP and federated learning, on recommendation performance, we compare with two variants of the model: *w/o LDP* and *Centralized. w/o LDP* refers that we directly transfer transformed user embedding across domains without LDP in P2FCDR. Besides, *Centralized* is a model directly trained with centralized data, which architecture is exactly the same as P2FCDR. We see its performance as the upper bound of the performance of P2FCDR. We compare the performance of these methods with ours on Movie & Book and present the metrics in Figure 2. The loss in recommendation performance is not significant, and we can con-



Figure 2: The performance impact of privacy-preserving technologies on Movie & Book dataset.



Figure 3: Impact of hyperparameter λ in LDP on recommendation performance on Movie & Book dataset.

clude that the P2FCDR provides measurable privacy protection with an acceptable loss of performance.

Privacy Budget

 λ is a hyperparameter in the LDP, which controls the strength of the Laplace noise and the privacy budget. To show the impact of the parameter λ , we evaluate our model under different λ in the range of [0, 0.1], step is 0.02. We show the results in terms of HR@10 and NDCG@10 for different λ on Movie & Book dataset in Figure 3. As λ increases, it indicates that the noise added to the transferred user embedding is larger, which reduces the privacy budget and data availability, resulting in a continuous decrease in recommendation performance. Performance of the model decreases significantly when λ is greater than 0.04, and we must have a trade-off between performance and privacy protection. Thus, we prefer to set λ to 0.02, which protects privacy without causing a prominent loss of performance.

Conclusion

In this paper, we propose a privacy-preserving CDR framework P2FCDR, which aims to satisfy more practical situation and enhance cross-domain information fusion at feature level. For this, we utilize federated learning mechanism to store user data locally instead of sharing it with other domains for protecting data privacy of business partners. Meanwhile, we derive a gated selecting vector based on the similarity between the in-domain and cross-domain user embedding for extracting useful signals that are highly related to the target domain. We empirically study the cross-domain recommendation performance of our proposed P2FCDR and analyze the effectiveness of privacy protection.

Acknowledgments

This work was supported by the National Key R&D Program of China (2019YFB2101805), the National Natural Science Foundation of China (72061127001), and the Beijing Natural Science Foundation (4212021).

References

Chai, D.; Wang, L.; Chen, K.; and Yang, Q. 2020. Secure federated matrix factorization. *IEEE Intelligent Systems*, 36(5): 11–20.

Chen, C.; Wu, H.; Su, J.; Lyu, L.; Zheng, X.; and Wang, L. 2022. Differential Private Knowledge Transfer for Privacy-Preserving Cross-Domain Recommendation. In *WWW*, 1455–1465.

Chen, X.; Zhang, Y.; Tsang, I. W.; Pan, Y.; and Su, J. 2020. Towards equivalent transformation of user preferences in cross domain recommendation. *TOIS*.

Choi, W.-S.; Tomei, M.; Vicarte, J. R. S.; Hanumolu, P. K.; and Kumar, R. 2018. Guaranteeing local differential privacy on ultra-low-power systems. In *ISCA*, 561–574. IEEE.

Gao, C.; Chen, X.; Feng, F.; Zhao, K.; He, X.; Li, Y.; and Jin, D. 2019a. Cross-domain recommendation without sharing user-relevant data. In *WWW*, 491–502.

Gao, C.; Huang, C.; Yu, Y.; Wang, H.; Li, Y.; and Jin, D. 2019b. Privacy-preserving cross-domain location recommendation. *IMWUT*, 3(1): 1–21.

Gao, C.; Li, Y.; Feng, F.; Chen, X.; Zhao, K.; He, X.; and Jin, D. 2021. Cross-domain recommendation with bridge-item embeddings. *TKDD*, 16(1): 1–23.

He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *WWW*, 173–182.

Hu, G.; Zhang, Y.; and Yang, Q. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *CIKM*, 667–676.

Krishnan, A.; Das, M.; Bendre, M.; Yang, H.; and Sundaram, H. 2020. Transfer learning via contextual invariants for one-to-many cross-domain recommendation. In *SIGIR*, 1081–1090.

Li, P.; and Tuzhilin, A. 2020. Ddtcdr: Deep dual transfer cross domain recommendation. In *WSDM*, 331–339.

Li, P.; and Tuzhilin, A. 2021. Dual metric learning for effective and efficient cross-domain recommendations. *TKDE*, (01): 1–1.

Liu, J.; Zhao, P.; Zhuang, F.; Liu, Y.; Sheng, V. S.; Xu, J.; Zhou, X.; and Xiong, H. 2020a. Exploiting aesthetic preference in deep cross networks for cross-domain recommendation. In *WWW*, 2768–2774.

Liu, M.; Li, J.; Li, G.; and Pan, P. 2020b. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *CIKM*, 885–894.

Liu, S.; Xu, S.; Yu, W.; Fu, Z.; Zhang, Y.; and Marian, A. 2021. FedCT: Federated collaborative transfer for recommendation. In *SIGIR*, 716–725.

Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019a. Neural graph collaborative filtering. In *SIGIR*, 165–174.

Wang, Y.; Feng, C.; Guo, C.; Chu, Y.; and Hwang, J.-N. 2019b. Solving the sparsity problem in recommendations via cross-domain item embedding based on co-clustering. In *WSDM*, 717–725.

Xie, R.; Ling, C.; Wang, Y.; Wang, R.; Xia, F.; and Lin, L. 2021. Deep feedback network for recommendation. In *IJ*-*CAI*, 2519–2525.

Xue, H.-J.; Dai, X.; Zhang, J.; Huang, S.; and Chen, J. 2017. Deep matrix factorization models for recommender systems. In *IJCAI*, 3203–3209.

Yan, D.; Zhao, Y.; Yang, Z.; Jin, Y.; and Zhang, Y. 2022. FedCDR: Privacy-preserving federated cross-domain recommendation. *Digital Communications and Networks*.

Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated machine learning: Concept and applications. *TIST*, 10(2): 1–19.

Zhang, Z.; Jin, X.; Li, L.; Ding, G.; and Yang, Q. 2016. Multi-domain active learning for recommendation. In *AAAI*, 2358–2364.

Zhao, C.; Li, C.; Xiao, R.; Deng, H.; and Sun, A. 2020. CATN: Cross-domain recommendation for cold-start users via aspect transfer network. In *SIGIR*, 229–238.

Zhu, F.; Chen, C.; Wang, Y.; Liu, G.; and Zheng, X. 2019. Dtcdr: A framework for dual-target cross-domain recommendation. In *CIKM*, 1533–1542.

Zhu, F.; Wang, Y.; Chen, C.; Liu, G.; and Zheng, X. 2020. A Graphical and Attentional Framework for Dual-Target Cross-Domain Recommendation. In *IJCAI*, 3001–3008.

Zhu, F.; Wang, Y.; Chen, C.; Zhou, J.; Li, L.; and Liu, G. 2021. Cross-domain recommendation: challenges, progress, and prospects. In *IJCAI*, 4721–4728.