

TransFR: Transferable Federated Recommendation with Adapter Tuning on Pre-trained Language Models

Honglei Zhang^{1,2}, Zhiwei Li³, Haoxuan Li⁴, Xin Zhou⁵, Jie Zhang⁵, Yidong Li^{1,2*}

¹Key Laboratory of Big Data & Artificial Intelligence in Transportation, Ministry of Education, China

²Beijing Jiaotong University, China

³University of Technology Sydney, Australia

⁴Peking University, China

⁵Nanyang Technological University, Singapore
{honglei.zhang, ydli}@bjtu.edu.cn

Abstract

Federated recommendations (FRs), facilitating multiple local clients to collectively learn a global model without disclosing user private data, have emerged as a prevalent on-device service. In conventional FRs, a dominant paradigm is to utilize discrete identities to represent clients and items, which are then mapped to domain-specific embeddings to participate in model training. Despite considerable performance, we reveal three inherent limitations that can not be ignored in federated settings, i.e., non-transferability across domains, ineffectiveness in cold-start settings, and potential privacy violations during federated training. To this end, we propose a transferable federated recommendation model, *TransFR*, which delicately incorporates the general capabilities empowered by pre-trained models and the personalized abilities by fine-tuning local private data. Specifically, it first learns domain-agnostic representations of items by exploiting pre-trained models with public textual corpora. To tailor for FR tasks, we further introduce efficient federated adapter-tuning and post-adaptation personalization, which facilitate personalized adapters for each client by fitting local private data. We theoretically prove the advantages of incorporating adapter tuning in FRs regarding both effectiveness and privacy. Through extensive experiments, we show that our *TransFR* surpasses state-of-the-art FRs on transferability.

Introduction

Federated recommendations (FRs), owing to the benefits of efficient and secure distributed learning (Feng et al. 2018), have recently attracted a lot of interest both in academia and industry (Zhang et al. 2025; Du et al. 2024; Wu et al. 2022). Empirically, it can be ingeniously generalized into a unified FR framework (Zhang et al. 2024b,e; Chen et al. 2026), which initially involves a user encoder and a fusion module in local client, and an item encoder in central server, as illustrated in Fig. 1. Many canonical works have expanded this research line in three directions. For instance, HPFL explores learning a powerful user encoder (Wu et al. 2021), RULE focuses on modeling efficient item embeddings (Chen et al. 2021), and Zhang *et al.* aims to design

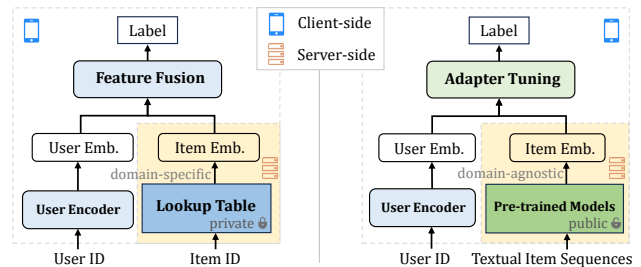


Figure 1: Illustration of the differences between standard FR (left) and our *TransFR* model (right). Existing work operates as a lookup table with domain-specific item embeddings, while *TransFR* exploits pre-trained language models to produce universal domain-agnostic embeddings.

adaptive fusion modules (Zhang et al. 2023a). Subsequent research has built upon this foundation by adding textual data into FR fields to mitigate the data sparsity issue (Li et al. 2025c; Hu et al. 2025). For instance, IFedRec incorporates the textual descriptions of items based on user interactions to improve model performance (Zhang et al. 2024a). It is worth noting that all these federated models mainly utilize discrete identities (IDs) to represent items, which have been an effective solution within single domains.

Although ID-based FR models have been well established, they routinely suffer from the following multilevel challenges in cross-domain tasks. First, it lacks *transferability* across platforms, as item IDs are generally not shared in practical scenarios. To be specific, the same IDs across different domains may have different meanings, hence the embeddings trained with these IDs are domain-specific, which seriously limits the development of transferable FR models (Muhammad et al. 2020). Second, it fails to handle the *cold-start issue* due to the ID-based item embeddings. Specifically, a well-trained embedding highly relies on sufficient ID interactions, which cannot provide satisfactory results when users and items have few interactions (Zhang et al. 2024a). Third, it poses a *privacy breach issue* in FR settings since the item embeddings are trained directly with

*Corresponding author.

user private data (Zhang et al. 2024d). In summary, none of these FR models take into account all the issues, which are the primary challenges for real-world FRs.

Given the challenges of existing work, we believe it is essential to develop a universal FR model that can effectively transfer to a new domain while ensuring unseen items can be handled well in cold-start settings. To achieve this, we can simply resort to the notion of retraining techniques to train the model from scratch with new domain-specific data (Zhang et al. 2020). Besides, we can augment side information or perturbations during the retraining process to alleviate the cold-start and privacy issues, respectively (Zhang et al. 2024a; Li, Song, and Fragouli 2020). However, such a straight solution brings a significant efficiency bottleneck since retraining the model involves multiple rounds of communication in federated learning (Guo et al. 2025; Feng et al. 2025), especially in resource-limited cross-device tasks. Moreover, retraining implies directly accessing to local raw data in the new domain, undoubtedly increasing the risk of user privacy leakage. Hence, it is imperative to design an efficient transferable FR framework, which can represent items in a domain-agnostic manner, following a privacy-preserving paradigm.

Inspired by recent findings suggesting that ID features of items can be replaced by textual features (Li et al. 2025a), we present *TransFR*, a principled transferable FR framework built on foundation models and federated adapter tuning techniques. Specifically, it represents items in a domain-agnostic way by purely utilizing fine-grained textual embeddings (as shown in the right panel of Fig. 1) instead of discrete item IDs. During this phase, language models like BERT are trained on publicly available text corpora on resource-rich servers. It can be understood that before the enactment of GDPR¹, these models could be trained with centrally collected data to achieve superior performance. After that, clients conduct efficient federated adapter tuning process, adapting it to the recommender task without transmitting raw data. This phase can be seen as post-GDPR period, adhering to legal restrictions on data stored at local clients. Finally, we introduce post-adaptation personalization to further enhance the personalization of FR tasks. Overall, our main contributions are listed as follows:

- We propose a unified FR model, *TransFR*, which can delicately exploits universal capabilities from foundation models and the personalized abilities from fine-tuning techniques. To the best of our knowledge, this is the first work focusing on textual transferability in FRs.
- To tailor for FRs, we introduce a simple yet effective federated adapter tuning to endow generic language models with collaborative signals. We further prove its effectiveness and privacy of adapter tuning in FRs theoretically, ensuring a dual guarantee of efficiency and privacy.
- To achieve personalized FRs, we introduce a post-adaptation personalization module, with locally fitting a prediction head on private data distributions, effectively mitigating data heterogeneity across clients.

¹<https://gdpr-info.eu/>

- Extensive experiments on several datasets show the advantages of our model on effectiveness and transferability over state-of-the-art counterparts.

Related Work

In this section, we review two relevant areas to this work, *i.e.*, federated recommendation (Zhang et al. 2025) and transferable recommendation (Zang et al. 2022).

Federated Recommendation

Federated recommendation (FR) has emerged as a promising mobile service in recent years, due to its ability to preserve user privacy (Li et al. 2025b; Yi et al. 2025; Chen et al. 2025). FCF (Ammad-Ud-Din et al. 2019) and FedRec (Lin et al. 2020) are two pioneering works to utilize the discrete identities (IDs) to learn user/item embeddings. In general, there are three major components in the standard FR framework, *i.e.*, user encoder, item encoder, and fusion module (Zhang et al. 2023a). As for the user encoder, Luo *et al.* proposed a personalized user encoder by combining the global and local clustered models (Luo, Xiao, and Song 2022). As for the item encoders, researchers have been working to address the privacy (Li, Song, and Fragouli 2020) and efficiency (Zhang et al. 2023b) issues associated with representing item embeddings with IDs. As for the fusion module, FedCIA proposed a similarity-based model fusion strategy to prevent information loss during the aggregation process (Han et al. 2025). Recent work, IFedRec (Zhang et al. 2024a), leveraged item attributes as side information to alleviate the cold-start issue in FR framework. Different from relying on coarse-grained IDs, our work represents items with fine-grained textual tokens to alleviate the transferability, cold-start and privacy issues of FRs.

Transferable Recommendation

Transferable recommendation aims to improve target-domain performance by leveraging auxiliary data from other domains, mitigating data sparsity issues (Zang et al. 2022; Bao et al. 2022; Zhang et al. 2024c). Early approaches relied on explicit or implicit domain overlaps, using mapping functions for knowledge transfer (Man et al. 2017). For instance, EMCDCR used a multi-layer perceptron for mapping to improve the robustness of target domain (Man et al. 2017). Distinct from the above work, subsequent efforts focus on learning universal representations with language models across different domains (Yuan et al. 2023). For example, UniSRec utilized the associated descriptions of items to learn transferable embeddings across different scenarios, showing the transferability in a cross-platform setting (Hou et al. 2022). Subsequent works aim to enhance transferability in federated settings. For example, FedDCSR adopted disentangled ID-based embeddings for transfer, but may expose sensitive data during embedding transmission (Zhang et al. 2024d). FedCLR inferred user interaction distributions in the source domain, reconstructs target-domain interactions via a generative model (Wang et al. 2025). Unlike prior methods that transfer item embeddings tied to private user data, our approach first learns transferable public parameters from open

corpora, and then applies adapter tuning with private data, ensuring both privacy and transferability.

The Proposed TransFR Framework

In this section, we introduce the details of TransFR model.

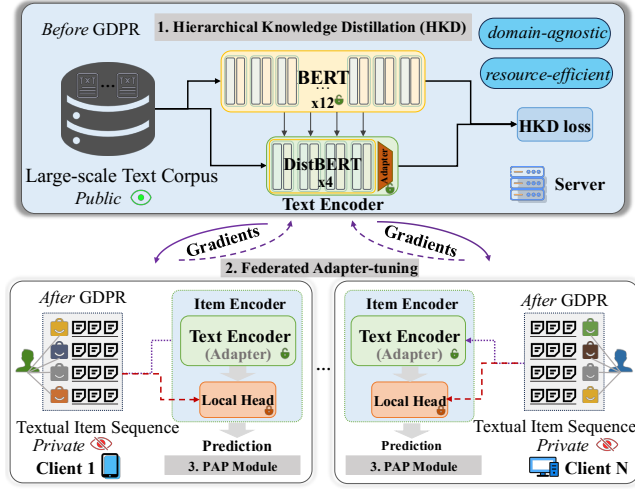


Figure 2: The schematic diagram of TransFR, which includes hierarchical knowledge distillation, federated adapter-tuning and post-adaptation personalization (PAP).

Preliminary

We consider a standard federated recommendation scenario with n users $\mathcal{U} = \{u\}$ and m items $\mathcal{I} = \{i\}$ on a central server. Each user u holds local feedback data \mathcal{D}_u consisting of interactions (u, i, y_{ui}) , where $y_{ui} = 1$ indicates user-item interaction. The goal is to predict \hat{y}_{ui} for unseen items and recommend the top K items to each user. In transferable recommendation, we define a source domain S and multiple target domains $T = \{T_1, \dots, T_D\}$, each with their own user/item sets. In the source domain (and similarly in target domains), there exists a user set \mathcal{U}^s (\mathcal{U}^t) and an item set \mathcal{I}^s (\mathcal{I}^t), comprising $|\mathcal{U}^s|$ ($|\mathcal{U}^t|$) users and $|\mathcal{I}^s|$ ($|\mathcal{I}^t|$) items. Our work focuses on cross-domain item recommendation for shared clients, utilizing domain-agnostic textual embeddings as an exploration in federated settings.

The TransFR Model

Overview Fig. 2 illustrates our framework. First, a central server pre-trains a large BERT model φ_{BERT} on public text data. To suit resource-limited devices, hierarchical knowledge distillation produces a smaller DistBERT $\varphi_{\text{DistBERT}}$. DistBERT is then combined with an Adapter ψ for efficient federated adapter-tuning, forming the Text Encoder $g = \psi \circ \varphi$. This encoder is deployed to user devices for local training. For personalization, a post-adaptation personalization module adds a local head h_u per client, creating a personalized Item Encoder $\phi_i = h_u \circ g$. For the user representation \mathbf{p}_u on local clients, we utilize local data to perform local training of the User Encoder ϕ_u . Ultimately, we

can utilize the fusion module $f(\cdot)$ to facilitate the prediction function for recommending item i to user u , denoted as $\hat{y}_{ui} = f(\mathbf{p}_u, \mathbf{q}_i)$, where \mathbf{q}_i denotes the produced embedding of item i . This process iterates until convergence.

Server-side Execution Procedure. In TransFR, the server leverages its computational resources to pre-train and distill a universal language model from centralized text corpora. Given the availability of mature pre-trained models, we skip pre-training and directly use BERT (Devlin et al. 2018) pre-trained on BooksCorpus (Zhu et al. 2015) and Wikipedia. These datasets predate GDPR, permitting centralized use. We then apply hierarchical knowledge distillation to produce an efficient model for resource-limited devices.

Hierarchical Knowledge Distillation. To enable efficient deployment on end-user devices, we introduce a hierarchical knowledge distillation (HKD) method tailored for BERT (Jiao et al. 2019). HKD transfers knowledge from a larger teacher BERT model with l Transformer layers to a smaller student model with s layers ($l > s$), both sharing the same architecture except for depth. The goal is to reduce model size and speed up inference while maintaining accuracy by minimizing a defined distillation loss.

$$\mathcal{L}_{\text{HKD}} = \sum_{x \in \mathcal{X}} \sum_{z=0}^z \mathcal{L}_{\text{layer}} \left(\varphi_{\text{DistBERT}}^z(x), \varphi_{\text{BERT}}^{p(z)}(x) \right), \quad (1)$$

where x is the textual input and \mathcal{X} the centralized training set collected before GDPR enforcement. The mapping function $l = p(z)$ maps multiple layers of the teacher model onto a single layer of the student model. Specifically, we set $p(z) = 3 \times z$, so DistBERT distills knowledge from every third teacher layer. The distillation loss targets individual BERT layers, primarily Transformer layers with multi-head attention. We design an attention-based loss that aligns the student’s attention matrices with the teacher’s, formally defined as follows:

$$\mathcal{L}_{\text{attn}} = \frac{1}{h} \sum_{i=1}^h (\mathbf{A}_i^{\text{stu}} - \mathbf{A}_i^{\text{tea}})^2, \quad (2)$$

where h is the number of attention heads, and $\mathbf{A}_i \in \mathbb{R}^{L \times L}$ is the attention matrix of the i -th head, with L as the input length. Inspired by studies showing BERT’s attention captures rich knowledge (Clark et al. 2019), we use mean squared error loss to transfer knowledge from teacher to student. Besides attention, we also distill knowledge from the embedding layer, the objective is defined as:

$$\mathcal{L}_{\text{embd}} = (\mathbf{E}^{\text{stu}} \mathbf{W} - \mathbf{E}^{\text{tea}})^2, \quad (3)$$

where the matrices $\mathbf{E}^{\text{stu}} \in \mathbb{R}^{L \times f}$ and $\mathbf{E}^{\text{tea}} \in \mathbb{R}^{L \times f'}$ correspond to the embeddings of student and teacher models, respectively. f and f' represent the embedding sizes for the student and teacher models, with f typically being smaller than f' to achieve a more compact student model. The matrix $\mathbf{W} \in \mathbb{R}^{f \times f'}$ is a learnable linear transformation aligning the dimensional space of the student network’s embeddings with that of the teacher network’s embeddings. Since

in the following sections we primarily utilize the distilled smaller student model $\varphi_{\text{DistBERT}}$, we will abusively use the symbol φ to denote $\varphi_{\text{DistBERT}}$ for clarity.

Efficient Federated Adapter-Tuning. To adapt this pre-trained model for recommendation tasks, it is necessary to collaboratively fine-tune the model between the server and clients using local private data. To achieve this goal, we propose an efficient federated adapter-tuning (FAT) method, which introduces a small, learnable Adapter on top of the frozen model, effectively balancing privacy and efficiency. Specifically, we incorporate an Adapter ψ_{Adapter} into the DistBERT model $\varphi_{\text{DistBERT}}$, resulting in a Text Encoder $g(\cdot)$ that explicitly combines both generality and personalization. The encoder $g(\cdot)$ can be formally represented as:

$$g(\cdot) = \underbrace{\psi_{\text{Adapter}}}_{\text{trainable}} \circ \underbrace{\varphi_{\text{DistBERT}}}_{\text{frozen}}, \quad (4)$$

where the Adapter consists of multi-layer perceptrons. After the server aggregates the gradients \mathbf{g}_θ of Adapter ψ in a weighted sum fashion uploaded from the sampled clients \mathcal{U}_p , the global aggregation and update of the Adapter ψ as:

$$\theta_{\text{adapter}} = \theta_{\text{adapter}} - \eta_s \cdot \left(\frac{1}{|\mathcal{U}_p|} \sum_{u \in \mathcal{U}_p} \mathbf{g}_\theta^u \right), \quad (5)$$

where η_s is the learning rate adopted in the server, and \mathbf{g}_θ^u represents the gradients of the Adapter uploaded from client u . This aggregation rule follows FedAvg (McMahan et al. 2017). We present more theoretical analysis on model effectiveness and privacy guarantee about our simple yet effective federated adapter tuning. The local optimization process will be detailed in the following section.

Client-side Execution Procedure. In TransFR, clients handle local optimization for federated adapter-tuning by updating user embeddings and uploading adapter gradients for global aggregation. They also perform post-adaptation personalization with local data to enhance personalization.

Efficient Federated Adapter-Tuning. The clients mainly handle the local optimization steps during the efficient federated adapter-tuning. Given the user embedding \mathbf{p}_u and item embedding \mathbf{q}_i of item i stored in client u , we adopt inner product to predict the preference score of user u to item i in the local device. Here, we utilize the well-designed binary cross-entropy (BCE) loss, which is as follows:

$$\mathcal{L}_{\text{local}} = - \sum_{(u,i) \in \mathcal{D}_u} y_{ui} \log \hat{y}_{ui} + (1 - y_{ui}) \log (1 - \hat{y}_{ui}), \quad (6)$$

where $\mathcal{D}_u = \mathcal{D}_u^+ \cup \mathcal{D}_u^-$ and \mathcal{D}_u^+ denotes the set of observed interactions, *i.e.*, $y_{ui} = 1$, and \mathcal{D}_u^- denotes the set of negative instances, *i.e.*, $y_{ui} = 0$, which can be uniformly sampled from unobserved interactions. For every observed interaction, there are N negative instances sampled for efficient training. By optimizing the BCE loss in the local client, we can update \mathbf{p}_u by stochastic gradient descent as follows:

$$\mathbf{p}_u = \mathbf{p}_u - \eta_c \cdot \frac{\partial \mathcal{L}_{\text{local}}}{\partial \mathbf{p}_u}, \quad (7)$$

where η_c is learning rate in clients and the specific gradient form of $\frac{\partial \mathcal{L}}{\partial \mathbf{p}_u}$ is determined based on the adopted user encoder ϕ_u . Besides, we calculate the gradients $\mathbf{g}_\theta = \{\Delta_{\theta_{\text{bert}}}, \Delta_{\theta_{\text{adapter}}}\}$ of the local client for the global parameters $\theta_{\text{enc}} = \{\theta_{\text{bert}}, \theta_{\text{adapter}}\}$ in Text Encoder $g(\cdot) = \psi \circ \varphi$, with the loss function $\mathcal{L}_{\text{local}}$, which is as follows:

$$\Delta_{\theta_{\text{bert}}} = \frac{\partial \mathcal{L}_{\text{local}}}{\partial \theta_{\text{bert}}}, \quad \Delta_{\theta_{\text{adapter}}} = \frac{\partial \mathcal{L}_{\text{local}}}{\partial \theta_{\text{adapter}}}, \quad (8)$$

where the parameters of the pre-trained BERT θ_{bert} are frozen, hence $\Delta_{\theta_{\text{bert}}} = 0$ and $\mathbf{g}_\theta = \Delta_{\theta_{\text{adapter}}}$. After the selected participating clients \mathcal{U}_p completes the gradient calculation of global parameters, then only upload \mathbf{g}_θ to the server for global aggregation procedure.

Post-Adaptation Personalization. Inspired by recent pioneering work, which indicates that the success of centralized machine learning is often attributable to the shared global feature representation, while the heterogeneity across clients or tasks predominantly manifests in the labels (Lecun, Bengio, and Hinton 2015). To mitigate this heterogeneity, we introduce post-adaptation personalization (PAP) mechanism with a local prediction head $h_u(\cdot)$ on top of the Text Encoder $g(\cdot)$ for each client, which involves local training on the fine-tuned Item Encoder $\phi_i = h_u \circ g$ using private local data. This mechanism is designed to further fit the local data distributions, thereby facilitating personalized FRs. Specifically, we update the local prediction head using the local loss function $\mathcal{L}_{\text{local}}$ in Eq.(6), detailed below:

$$\theta_{\text{head}} = \theta_{\text{head}} - \eta_c \cdot \frac{\partial \mathcal{L}_{\text{local}}}{\partial \theta_{\text{head}}}, \quad (9)$$

where θ_{head} represents the trainable parameters of the local prediction head $h_u(\cdot)$, which is composed of multi-layer perceptrons. To reduce the risk of overfitting, the local head can adopt the same model structure as the Adapter, thereby reducing the learnable parameters in the Item Encoder. Intuitively, adding a personalized prediction head to each common pre-trained model can expand the optimization space of the training parameters, thereby yielding improved personalized prediction performance.

From a theoretical standpoint, this PAP strategy can be explained using domain adaptation theory (Ben-David et al. 2010). Let P_S denote the global source distribution (on which the shared encoder g is trained), and P_T denote the client-specific target distribution. Then, the generalization error on the client domain can be bounded by:

$$\epsilon_T(h_u \circ g) \leq \epsilon_S(g) + \frac{1}{2} \mathcal{D}_{\mathcal{H}}(P_S, P_T) + \lambda^*, \quad (10)$$

where ϵ_T and ϵ_S denote the expected prediction error on the target and source domains respectively, $\mathcal{D}_{\mathcal{H}}$ measures the divergence between the two distributions in hypothesis space \mathcal{H} , and λ^* is the irreducible error of the optimal joint hypothesis. By updating h_u locally to minimize ϵ_T , the PAP mechanism reduces the distribution gap $\mathcal{D}_{\mathcal{H}}$ and improves the personalization capacity of the model.

Theoretical Analysis

Next, we present theoretical analysis about model effectiveness and privacy guarantee by federated adapter tuning.

Model Effectiveness. In FRs, all clients collaboratively learn a shared item representation model φ by minimizing:

$$\min_{\varphi} \frac{1}{n} \sum_{u=1}^n \ell_u(\varphi)$$

where ℓ_u is the local loss for the u -th client. However, this global optimization may lead to sub-optimal performance in heterogeneous settings due to client distribution shifts.

We adopt a personalized formulation where clients jointly learn a shared item representation model φ , while locally adapting their own adapters ψ_u , the revised objective is:

$$\min_{\psi_u, \varphi} \frac{1}{n} \sum_{u=1}^n \min_{\psi_u \in \mathcal{H}} \ell_u(\psi_u \circ \varphi)$$

This scheme allows better personalization and representation alignment for downstream recommendation. Under a linearized setting, we provide the following result:

Theorem 1 (Effectiveness of Personalized Adapters). *Let the global item representation model be a linear transformation $\mathbf{A} \in R^{L \times f}$ and the local adapters for client u be $\mathbf{B}_u \in R^{f \times d}$. Then, we have:*

- The optimal solution to the personalized formulation

$$\min_{\mathbf{A}, \{\mathbf{B}_u\}} \frac{1}{n} \sum_{u=1}^n \|\mathbf{A}\mathbf{B}_u - \mathbf{A}^*\mathbf{B}_u^*\|_2^2$$

achieves zero loss when \mathbf{A} spans the same space as \mathbf{A}^* and $\{\mathbf{B}_u^*\}$ spans R^d .

- The standard formulation with a shared global model

$$\min_{\mathbf{A}, \mathbf{B}} \frac{1}{n} \sum_{u=1}^n \|\mathbf{A}\mathbf{B} - \mathbf{A}^*\mathbf{B}_u^*\|_2^2$$

incurs residual error increasing with the heterogeneity of $\{\mathbf{B}_u^*\}$.

Privacy Guarantee. We analyze the privacy benefits of our proposed federated adapter tuning in FR settings through an information-theoretic lens. Our key findings are summarized as follows.

Theorem 2 (Mutual Information Upper Bound). *Let ψ_u be the adapter parameters trained on user data \mathcal{D}_u , with bounded value range $[-C, C]$ and quantization resolution δ . The information leakage $I(\psi_u; \mathcal{D}_u)$ is upper bounded by:*

$$I(\psi_u; \mathcal{D}_u) \leq fd \cdot \log \left(\frac{2C}{\delta} \right)$$

where d is the bottleneck dimension in the adapter.

Theorem 3 (Mutual Information Bound from DP Mechanism). *If Gaussian noise is added to adapter parameters $\tilde{\psi}_u$ such that the update satisfies (ϵ, δ) -differential privacy (DP), then the mutual information leakage is bounded by:*

$$I(\tilde{\psi}_u; \mathcal{D}_u) \leq \epsilon^2 + \delta$$

These results suggest that federated adapter tuning leaks significantly less information compared to full-model or embedding updates, due to its structural bottleneck and optional DP perturbation.

Experiments

In this section, we conduct experiments to answer the following research questions:

RQ1: Does our proposed TransFR method outperform the state-of-the-art methods regarding transferability?

RQ2: Does our proposed TransFR surpass existing transferable baselines in cold-start settings?

RQ3: How do the specific components in our proposed method work for FR tasks, *i.e.*, pre-training, federated adapter-tuning, and post-adaptation personalization?

RQ4: How do certain hyperparameters within our method impact the FR tasks, *i.e.*, the number of layers in the adapter, the number of epochs in post-adaptation personalization?

Experimental Settings

Datasets & Metrics To benchmark our model against baselines, we perform experiments with the *Amazon* (Ni, Li, and McAuley 2019) dataset. We selected three categories as distinct domains: Movies, Music, and Books. Within each domain, we excluded interaction records lacking titles of items and applied filters to retain users with at least 20 interactions and items with a minimum of 30 interactions, in alignment with previous studies (He et al. 2017; Zhao et al. 2020). To assess our model and baselines, we employ a leave-one-out evaluation approach and two prevalent evaluation metrics in ranking tasks: HR and NDCG.

Comparison Models We adopt two kinds of benchmarks for overall comparisons, *i.e.*, centralized models, *i.e.*, CMF (Singh and Gordon 2008), EMCDR (Man et al. 2017), DisenCDR (Cao et al. 2022), and recent federated approaches, *i.e.*, FCF (Ammad-Ud-Din et al. 2019), IFedRec (Zhang et al. 2024a), FELLRec (Zhao et al. 2025), FedCT (Liu et al. 2021) and FedDCSR (Zhang et al. 2024d).

Experimental Results

In this section, we detail the experimental results on Amazon dataset to verify the effectiveness of the proposed model.

Overall Performance We compare centralized and federated models, with the first three methods being centralized and the last five federated. Most rely on ID-based item embeddings and perform knowledge transfer via explicit ID mapping. To address RQ1, we evaluate the transferability of our method and baselines, as shown in Table 1 and Table 2, reporting HR@10 and NDCG@10 on Amazon dataset. Transferability is measured by $\Delta = \frac{\text{Source} - \text{Target}}{\text{Source}}$, using HR@10 as an example. The results reveal several key insights to generalize across domains.

Centralized approaches can achieve better transferability, but they inevitably pose privacy risks. Among federated baselines, FedDCSR and FedCT outperforms FCF, IFedRec and FELLREC. This advantage stems from its use of explicit mapping to extract domain-relevant features. In contrast, FCF, IFedRec and FELLRec directly apply source-trained models to the target domain. IFedRec performs better than FCF due to its textual side information, while FCF relies on matrix factorization with sparse interaction data.

Models	Music→Books			Music→Movies			Books→Movies			Average		
	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓
CMF	0.2232	0.1833	17.88%	0.2245	0.1378	38.62%	0.1792	0.1455	18.81%	0.2090	0.1555	25.57%
EMCDR	0.2232	0.2149	3.72%	0.2245	0.1893	15.68%	0.1792	0.1652	7.81%	0.2090	0.1898	9.17%
DisenCDR	0.2404	0.2243	6.70%	0.2302	0.1988	13.61%	0.2105	0.1912	9.17%	0.2270	0.2048	9.49%
FCF	0.1533	0.1380	9.98%	0.1803	0.1154	36.00%	0.1220	0.1004	17.70%	0.1519	0.1179	22.34%
IFedRec	0.2193	0.1981	9.66%	0.2095	0.1810	13.61%	0.1912	0.1703	10.95%	0.2067	0.1831	11.40%
FELLRec	0.2215	0.2003	9.56%	0.2121	0.1844	13.06%	0.1939	0.1731	10.76%	0.2092	0.1859	11.13%
FedCT	0.2240	0.2034	9.20%	0.2142	0.1765	17.60%	0.1963	0.1585	19.24%	0.2115	0.1795	15.35%
FedDCSR	0.2417	0.2217	8.27%	0.2221	0.1881	15.31%	0.1994	0.1719	13.79%	0.2211	0.1939	12.29%
TransFR	0.2386	0.2335	2.14%	0.2311	0.2123	8.14%	0.2146	0.2089	2.66%	0.2281	0.2182	4.33%

Table 1: The results of our TransFR and the baselines in terms of transferability towards HR@10 on Amazon dataset. Δ represents the transferability capability from the source domain to the target domain. The best results in federated settings are in bold. \uparrow indicates that higher values correspond to better performance while \downarrow denotes the lower values, the better performance.

Models	Music→Books			Music→Movies			Books→Movies			Average		
	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓	Source↑	Target↑	Δ↓
CMF	0.1392	0.0955	31.39%	0.1132	0.0764	32.51%	0.0921	0.0662	28.12%	0.1148	0.0794	30.89%
EMCDR	0.1392	0.1121	19.47%	0.1132	0.0901	20.41%	0.0921	0.0705	23.45%	0.1148	0.0909	20.84%
DisenCDR	0.1403	0.1182	15.75%	0.1121	0.0923	17.66%	0.0995	0.0821	17.48%	0.1173	0.0975	16.87%
FCF	0.0982	0.0663	32.48%	0.0721	0.0583	19.14%	0.0798	0.0421	47.24%	0.0834	0.0556	33.35%
IFedRec	0.1010	0.0834	17.43%	0.0954	0.0673	29.45%	0.0831	0.0497	40.19%	0.0932	0.0668	28.30%
FELLRec	0.1145	0.0995	13.10%	0.1068	0.0812	23.97%	0.1033	0.0831	19.55%	0.1082	0.0879	18.76%
FedCT	0.1221	0.1122	8.11%	0.1144	0.0899	21.41%	0.1047	0.0845	19.29%	0.1137	0.0955	16.01%
FedDCSR	0.1192	0.1085	8.98%	0.1143	0.0892	21.96%	0.0943	0.0732	22.38%	0.1093	0.0903	17.36%
TransFR	0.1209	0.1133	6.29%	0.1132	0.0934	17.49%	0.1054	0.0893	15.28%	0.1132	0.0987	12.81%

Table 2: The results of our TransFR and the baselines in terms of transferability towards NDCG@10 on Amazon dataset.

Our TransFR method achieves the best target domain performance, with the lowest Δ values in HR@10 (4.33%) and NDCG@10 (12.81%), demonstrating superior transferability. Although it slightly underperforms FedCT and FedDCSR in some source domain cases, it consistently outperforms them in target domain metrics and transfer loss. This advantage arises from using universal textual representations that reduce semantic gaps across domains, unlike ID-based methods relying on explicit ID mappings.

Table 3 aims to address RQ2, reporting HR@10 and NDCG@10 on Amazon dataset in cold-start settings, where models are evaluated on unseen items in the target domain. We derive several key observations. First, the federated baseline FCF struggles due to its ID-based item representation. IFedRec offers slight improvement by adding textual data but still relies on IDs, limiting its effectiveness. FedDCSR uses disentangled representations to transfer knowledge across domains, partially alleviating cold-start issues. Our method surpasses both ID-based and ID-mapping baselines by leveraging domain-agnostic textual embeddings and federated adapter tuning. Unlike coarse-grained ID embeddings, our textual representations capture item semantics more effectively and remain consistent across domains.




Figure 3: Ablation study results towards HR@10 in the three transferable tasks on Amazon dataset.

Ablation Study To address RQ3, we conduct an ablation study to assess the contribution of each component, i.e., pre-training (PT), federated adapter-tuning (FAT), and post-adaptation personalization (PAP), as shown in Fig.3 and Fig.4. We compare three variants, each removing one module: “w/o PT” replaces pre-training with random initialization, “w/o FAT” omits federated adapter tuning, and “w/o PAP” removes local head adaptation. As shown in Fig.3 and Fig.4, we observe the following: (1) Removing pre-training significantly degrades performance, aligning with prior findings (Chen et al. 2022); (2) Federated adapter-tuning yields moderate gains, validating the effectiveness of our Adapter

Models	Music→Books		Music→Movies		Books→Movies	
	H@10 ↑	N@10 ↑	H@10 ↑	N@10 ↑	H@10 ↑	N@10 ↑
CMF	0.1833	0.0955	0.1378	0.0764	0.1455	0.0662
EMCDR	0.2149	0.1121	0.1893	0.0901	0.1652	0.0705
DisenCDR	0.2287	0.1123	0.1934	0.0921	0.1892	0.0764
FCF	0.1380	0.0663	0.1154	0.0583	0.1004	0.0421
IFedRec	0.1944	0.0982	0.1654	0.0631	0.1583	0.0681
FELLRec	0.2093	0.1052	0.1723	0.0845	0.1639	0.0701
FedCT	0.2138	0.1062	0.1802	0.0858	0.1708	0.0726
FedDCSR	0.2217	0.1085	0.1881	0.0892	0.1719	0.0732
TransFR	0.2310	0.1126	0.2001	0.0934	0.1932	0.0830

Table 3: Comparison results of TransFR and the baselines on Amazon dataset in cold-start settings.




Figure 4: Ablation study results towards NDCG@10 in the three transferable tasks on Amazon dataset.

module; (3) Post-adaptation personalization further boosts transferability, highlighting the role of the local head in personalized recommendation.

δ	Music→Books		Music→Movies		Books→Movies	
	H@10 ↑	N@10 ↑	H@10 ↑	N@10 ↑	H@10 ↑	N@10 ↑
0.0	0.2310	0.1126	0.2001	0.0934	0.1932	0.0830
0.1	0.2301	0.1123	0.1989	0.0928	0.1913	0.0802
0.2	0.2289	0.1101	0.1964	0.0903	0.1889	0.0775
0.3	0.2234	0.1089	0.1903	0.0875	0.1845	0.0734

Table 4: Results of TransFR with differential privacy.

We further incorporate the differential privacy to enhance privacy capability, with δ denoting noise intensity into adapters, which leads to only a slight performance decline, as shown in Table 4. This shows that TransFR maintains competitive performance while keeping privacy protection.

We evaluate the impact of different pre-trained models by comparing DistBERT and Llama-2-7B on Amazon dataset. As shown in Table 5, Llama-2-7B yields slightly better accuracy (+2.70% H@10, +2.73% N@10) than DistBERT, but at a high memory and communication overheads rise by 20,982%, and inference cost increases by 400%. These results suggest: (1) Large language models offer marginal accuracy gains but are inefficient for item embedding tasks. This finding is consistent with recent pioneer works (Fournier, Dupoux, and Dunbar 2020; Freestone and Santu 2024); (2) DistBERT provides a better trade-off be-

Models	Accuracy		Memory	Comm.	Inference
	H@10 ↑	N@10 ↑	MB ↓	MB ↓	s ↓
DistBERT	0.2182	0.0987	68	68	0.003
Llama-2-7B	0.2241	0.1014	14,336	14,336	0.015
<i>Gain</i>	2.70%	2.73%	-20,982%	-20,982%	-400%

Table 5: Ablation study comparison of pre-trained modules between DistBERT and Llama-2 in terms of recommendation accuracy, memory overhead (MB), communication (commu.) overhead (MB), and inference time (s).




Figure 5: Performance analysis of TransFR with varying numbers of layers in the adapter module.




Figure 6: Performance analysis of TransFR with different local epochs in post-adaptation personalization module.

tween performance and efficiency for cross-device tasks.

Sensitivity Analysis To address RQ4, we analyze results with different hyperparameters. As shown in Fig. 5, increasing the number of hidden layers generally improves performance in the target domain, though some tasks (e.g., Music → Books, Books → Movies) show diminishing returns after a certain point. Thus, a moderate increase in hidden layers can enhance cross-device recommendation performance. We further examine the effect of local training epochs in post-adaptation personalization on Amazon dataset, as shown in Fig.6. Performance generally improves with more epochs, but overfitting may occur in scenarios with limited local data. Balancing effectiveness and efficiency, we set the number of local training epochs to 6 for all client devices.

Conclusion

In this work, we propose a transferable FR model with adapter tuning on pre-trained language models. By replacing domain-specific item embedding tables with domain-agnostic pre-trained textual representations and adapter tuning, we address the transferability, cold-start, and privacy issues. Besides, we discuss the superiority of TransFR on several aspects from theoretical and experimental perspectives. In future work, we plan to incorporate more powerful multi-modal foundation models to further enhance model performance within the FR framework.

Acknowledgments

This work was supported by the Beijing Natural Science Foundation 4264109, in part by the National Science Foundation of China under Grant U2268203, 62402031, 62203040, 62306027 and 62372436, and in part by the Beijing Nova Program 20240484620.

References

- Ammad-Ud-Din, M.; Ivannikova, E.; Khan, S. A.; Oyomno, W.; Fu, Q.; Tan, K. E.; and Flanagan, A. 2019. Federated collaborative filtering for privacy-preserving personalized recommendation system. *arXiv preprint arXiv:1901.09888*.
- Bao, S.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2022. The minority matters: A diversity-promoting collaborative metric learning algorithm. *NeurIPS*, 2451–2464.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79(1): 151–175.
- Cao, J.; Lin, X.; Cong, X.; Ya, J.; Liu, T.; and Wang, B. 2022. Disencdr: Learning disentangled representations for cross-domain recommendation. In *SIGIR*, 267–277.
- Chen, H.-Y.; Tu, C.-H.; Li, Z.; Shen, H. W.; and Chao, W.-L. 2022. On the importance and applicability of pre-training for federated learning. In *ICLR*.
- Chen, J.; Zhang, H.; Li, H.; Zhang, C.; Li, Z.; and Li, Y. 2025. Beyond personalization: Federated recommendation with calibration via low-rank decomposition. *arXiv preprint arXiv:2506.09525*.
- Chen, J.; Zhang, H.; Zhang, C.; Luo, F.; and Li, Y. 2026. Breaking the aggregation bottleneck in federated recommendation: A personalized model merging approach. In *AAAI*.
- Chen, T.; Yin, H.; Zheng, Y.; Huang, Z.; Wang, Y.; and Wang, M. 2021. Learning elastic embeddings for customizing on-device recommenders. In *SIGKDD*, 138–147.
- Clark, K.; Khandelwal, U.; Levy, O.; and Manning, C. D. 2019. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Du, Y.; Luo, D.; Yan, R.; Wang, X.; Liu, H.; Zhu, H.; Song, Y.; and Zhang, J. 2024. Enhancing job recommendation through llm-based generative adversarial networks. In *AAAI*, 8363–8371.
- Feng, J.; Lai, Y.; Sun, H.; and Ren, B. 2025. Sadba: Self-adaptive distributed backdoor attack against federated learning. In *AAAI*, 16568–16576.
- Feng, J.; Yang, L. T.; Zhu, Q.; and Choo, K.-K. R. 2018. Privacy-preserving tensor decomposition over encrypted data in a federated cloud environment. *IEEE Trans. Dependable Secure Comput.*, 17(4): 857–868.
- Fournier, L.; Dupoux, E.; and Dunbar, E. 2020. Analogies minus analogy test: measuring regularities in word embeddings. In *CoNLL*, 365–375.
- Freestone, M.; and Santu, S. K. K. 2024. Word embeddings revisited: Do llms offer something new? *arXiv preprint arXiv:2402.11094*.
- Guo, X.; Yu, K.; Cui, L.; Yu, H.; and Li, X. 2025. Federated causally invariant feature learning. In *AAAI*, 16978–16986.
- Han, M.; Li, D.; Xia, J.; Liu, J.; Gu, H.; Zhang, P.; Gu, N.; and Lu, T. 2025. FedCIA: Federated collaborative information aggregation for privacy-preserving recommendation. In *SIGIR*, 1687–1696.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *The Web Conference*, 173–182.
- Hou, Y.; Mu, S.; Zhao, W. X.; Li, Y.; Ding, B.; and Wen, J.-R. 2022. Towards universal sequence representation learning for recommender systems. In *SIGKDD*, 585–593.
- Hu, J.; Hooi, B.; He, B.; and Wei, Y. 2025. Modality-independent graph neural networks with global transformers for multimodal recommendation. In *AAAI*, 11790–11798.
- Jiao, X.; Yin, Y.; Shang, L.; Jiang, X.; Chen, X.; Li, L.; Wang, F.; and Liu, Q. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436–444.
- Li, G.; Jing, L.; Wu, J.; Li, X.; Zhu, K.; and He, Y. 2025a. From ID-based to ID-free: Rethinking ID effectiveness in multimodal collaborative filtering recommendation. In *ACM MM*.
- Li, T.; Song, L.; and Fragouli, C. 2020. Federated recommendation system via differential privacy. In *ISIT*, 2592–2597.
- Li, Z.; Long, G.; Zhang, C.; Zhang, H.; Jiang, J.; and Zhang, C. 2025b. Learn to preserve personality: Federated foundation models in recommendations. *arXiv preprint arXiv:2506.11563*.
- Li, Z.; Long, G.; Zhou, T.; Jiang, J.; and Zhang, C. 2025c. Personalized federated collaborative filtering: A variational autoencoder approach. In *AAAI*, 18602–18610.
- Lin, G.; Liang, F.; Pan, W.; and Ming, Z. 2020. Fedrec: Federated recommendation with explicit feedback. *IEEE Intell. Syst.*, 36(5): 21–30.
- Liu, S.; Xu, S.; Yu, W.; Fu, Z.; Zhang, Y.; and Marian, A. 2021. FedCT: Federated collaborative transfer for recommendation. In *SIGIR*, 716–725.
- Luo, S.; Xiao, Y.; and Song, L. 2022. Personalized federated recommendation via joint representation learning, user clustering, and model adaptation. In *CIKM*, 4289–4293.
- Man, T.; Shen, H.; Jin, X.; and Cheng, X. 2017. Cross-domain recommendation: An embedding and mapping approach. In *IJCAI*, 2464–2470.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTAT*, 1273–1282.

- Muhammad, K.; Wang, Q.; O'Reilly-Morgan, D.; Tragos, E.; Smyth, B.; Hurley, N.; Geraci, J.; and Lawlor, A. 2020. Fedfast: Going beyond average for faster training of federated recommender systems. In *SIGKDD*, 1234–1242.
- Ni, J.; Li, J.; and McAuley, J. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *EMNLP*, 188–197.
- Singh, A. P.; and Gordon, G. J. 2008. Relational learning via collective matrix factorization. In *SIGKDD*, 650–658.
- Wang, Q.; Zhao, Y.; Zhang, Y.; Zhang, Y.; Deng, S.; and Yang, Y. 2025. Federated contrastive learning for cross-domain recommendation. *IEEE Trans. Serv. Comput.*
- Wu, J.; Liu, Q.; Huang, Z.; Ning, Y.; Wang, H.; Chen, E.; Yi, J.; and Zhou, B. 2021. Hierarchical personalized federated learning for user modeling. In *The Web Conference*, 957–968.
- Wu, J.; Qian, B.; Li, Y.; Gao, Z.; Ju, M.; Yang, Y.; Zheng, Y.; Gong, T.; Li, C.; and Zhang, X. 2022. Leveraging multiple types of domain knowledge for safe and effective drug recommendation. In *CIKM*, 2169–2178.
- Yi, L.; Yu, H.; Ren, C.; Wang, G.; Liu, X.; and Li, X. 2025. pFedES: Generalized proxy feature extractor sharing for model heterogeneous personalized federated learning. In *AAAI*, 22146–22154.
- Yuan, Z.; Yuan, F.; Song, Y.; Li, Y.; Fu, J.; Yang, F.; Pan, Y.; and Ni, Y. 2023. Where to go next for recommender systems? id-vs. modality-based recommender models revisited. In *SIGIR*, 2639–2649.
- Zang, T.; Zhu, Y.; Liu, H.; Zhang, R.; and Yu, J. 2022. A survey on cross-domain recommendation: taxonomies, methods, and future directions. *ACM Trans. Inf. Syst.*, 41(2): 1–39.
- Zhang, C.; Long, G.; Zhang, Z.; Li, Z.; Zhang, H.; Yang, Q.; and Yang, B. 2025. Personalized recommendation models in federated settings: A survey. *IEEE Trans. Knowl. Data Eng.*
- Zhang, C.; Long, G.; Zhou, T.; Yan, P.; Zhang, Z.; and Yang, B. 2023a. Dual personalization on federated recommendation. In *IJCAI*, 4558–4566.
- Zhang, C.; Long, G.; Zhou, T.; Zhang, Z.; Yan, P.; and Yang, B. 2024a. When federated recommendation meets cold-start problem: Separating item attributes and user interactions. In *The Web Conference*, 3632–3642.
- Zhang, H.; Li, H.; Chen, J.; Cui, S.; Yan, K.; Wuerkaixi, A.; Zhou, X.; Shen, Z.; and Li, Y. 2024b. Beyond similarity: Personalized federated recommendation with composite aggregation. *arXiv preprint arXiv:2406.03933*.
- Zhang, H.; Luo, F.; Wu, J.; He, X.; and Li, Y. 2023b. LightFR: Lightweight federated recommendation with privacy-preserving matrix factorization. *ACM Trans. Inf. Syst.*, 41(4): 1–28.
- Zhang, H.; Wang, S.; Li, H.; Zheng, C.; Chen, X.; Liu, L.; Luo, S.; and Wu, P. 2024c. Uncovering the propensity identification problem in debiased recommendations. In *ICDE*, 653–666.
- Zhang, H.; Zheng, D.; Yang, X.; Feng, J.; and Liao, Q. 2024d. FedDCSR: Federated cross-domain sequential recommendation via disentangled representation learning. In *SDM*, 535–543.
- Zhang, H.; Zhou, X.; Shen, Z.; and Li, Y. 2024e. PrivFR: Privacy-enhanced federated recommendation with shared hash embedding. *IEEE Trans. Neural Networks Learn. Syst.*, 36(1): 32–46.
- Zhang, Y.; Feng, F.; Wang, C.; He, X.; Wang, M.; Li, Y.; and Zhang, Y. 2020. How to retrain recommender system? A sequential meta-learning method. In *SIGIR*, 1479–1488.
- 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.
- Zhao, J.; Wang, W.; Xu, C.; Ng, S.-K.; and Chua, T.-S. 2025. A federated framework for llm-based recommendation. In *Findings of NAACL*.
- Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV*, 19–27.