

HeteFedRec: Federated Recommender Systems with Model Heterogeneity

Wei Yuan¹ Liang Qu¹ Lizhen Cui² Yongxin Tong³ Xiaofang Zhou⁴ Hongzhi Yin^{1*}

¹The University of Queensland, ²Shandong University, ³Beihang University

⁴The Hong Kong University of Science and Technology

¹{w.yuan, l.qu, h.yin1}@uq.edu.au, ²clz@sdu.edu.cn

³yxtong@buaa.edu.cn, ⁴zxf@cse.ust.hk

Abstract—Owing to the nature of privacy protection, federated recommender systems (FedRecs) have garnered increasing interest in the realm of on-device recommender systems. However, most existing FedRecs only allow participating clients to collaboratively train a recommendation model of the same public parameter size. Training a model of the same size for all clients can lead to suboptimal performance since clients possess varying resources. For example, clients with limited training data may prefer to train a smaller recommendation model to avoid excessive data consumption, while clients with sufficient data would benefit from a larger model to achieve higher recommendation accuracy. To address the above challenge, this paper introduces HeteFedRec, a novel FedRec framework that enables the assignment of personalized model sizes to participants. Specifically, we present a heterogeneous recommendation model aggregation strategy, including a unified dual-task learning mechanism and a dimensional decorrelation regularization, to allow knowledge aggregation among recommender models of different sizes. Additionally, a relation-based ensemble knowledge distillation method is proposed to effectively distill knowledge from heterogeneous item embeddings. Extensive experiments conducted on three real-world recommendation datasets demonstrate the effectiveness and efficiency of HeteFedRec in training federated recommender systems under heterogeneous settings.

I. INTRODUCTION

Recently, there has been a significant surge in demand for recommender systems in various online services [1]–[7], since they can effectively alleviate information overload by actively discovering users’ potential interests. Conventional recommender systems are trained in a central server with collected user private data, which takes risks of data leakage and raises privacy concerns [8]. With the growing awareness of user privacy protection and some recently released regulations, such as GDPR¹ in the EU, CCPA² in the USA, and PIPL³ in China, it is becoming harder and even infeasible for online platforms to learn recommendation models on centrally collected user data [9].

Federated learning [10], as a privacy-preserving paradigm, has achieved remarkable success across various scenarios [11]. In order to tackle privacy concerns, researchers have sought to harness federated learning for the development of privacy-aware recommender systems namely federated recommender




Fig. 1: The distribution of users’ item interaction numbers. systems (FedRecs), where clients⁴ can collaboratively train a global recommendation model on their local devices without sharing their personal data with any other participants.

Ammand et al. [12] proposed the pioneering federated recommendation framework. In this framework, a recommender system is divided into public and private parameters. Users and a central server collaboratively learn a recommendation model by transmitting and aggregating the public parameters using specific aggregation strategies. Due to the advantage of privacy protection, several extended versions [9], [13]–[16] have been developed based on this fundamental framework to enhance the performance of FedRec.

Although many achievements have been made recently, almost all existing FedRecs operate under the setting where clients collaboratively learn a global recommendation model of the same size since their aggregation strategies (e.g., FedAvg [10]) can only process homogeneous public parameters, as shown in Fig. 2. This setting conflicts with the real applications where clients often possess heterogeneous resources, such as varying amounts of training data⁵. Fig. 1 depicts the user-item interaction number distribution on three recommendation datasets derived from real-world platforms. The standard deviation of interaction numbers is about 154.2, 79.8, and 105.2 for MovieLens-1M (ML) [17], Anime, and Douban datasets, while the average of interaction amounts is 132.8, 96.1, and 143.7, respectively, indicating the substantial difference of data size among clients. Under this circumstance, training a global recommendation model of the same size on all clients would fall short of achieving the true global

⁴In this paper, the terms “client” and “user” are used interchangeably, as each client is solely responsible for one user to ensure privacy protection.

⁵While this paper emphasizes the necessity of model heterogeneity using data size variance as an example, it’s crucial to recognize that such heterogeneity can also tackle some other resource diversity problems, including disparities in computational power, energy constraints, bandwidth, and more.

* Corresponding author.

¹<https://gdpr-info.eu>

²<https://oag.ca.gov/privacy/ccpa>

³<https://personalinformationprotectionlaw.com>

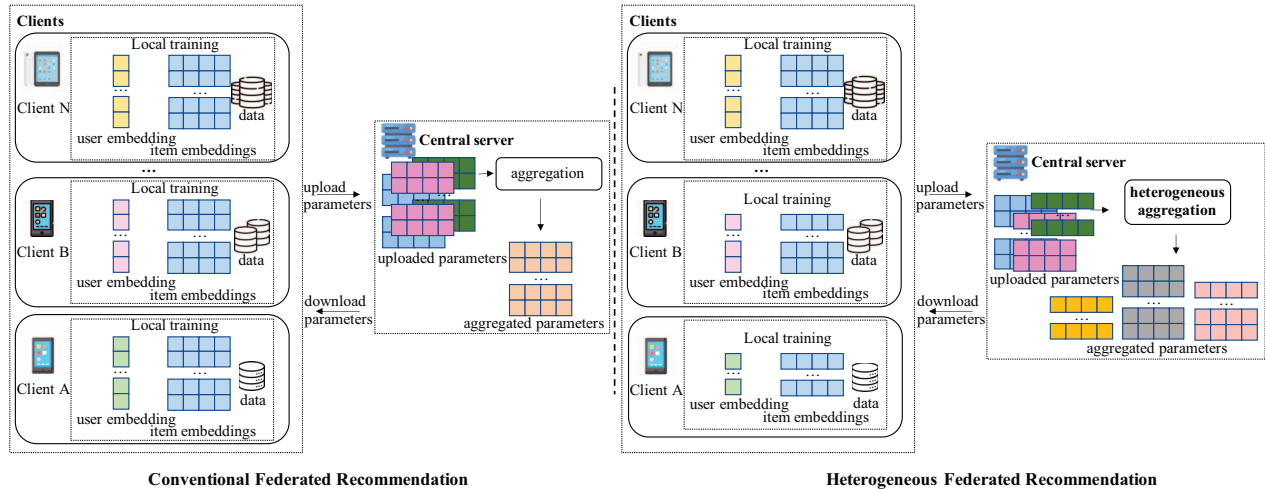


Fig. 2: Conventional Federated Recommender Systems v.s. Heterogeneous Federated Recommender Systems. Heterogeneous federated recommender systems assign public parameters with customized sizes considering clients' data amounts.

optimum [18]. Clients with limited data would struggle to support the training of the large global model, and their model updates might even have a negative impact on the performance of the global model, as their local models can only be updated for few times by gradient decent [19] on insufficient data. Conversely, clients with abundant data would prefer to train a larger model to obtain more appropriate recommendations. Given these limitations, there is an urgent need for a federated recommendation framework that enables the collaborative training of heterogeneous models⁶, tailored to the specific resource scale available to each client.

In federated learning, several endeavours have been undertaken to achieve model heterogeneity [20]. These methods can generally be classified into two categories. The first [21]–[23] involves leveraging knowledge distillation [24] to exchange knowledge over heterogeneous models via a public reference dataset. Another solution entails designing new aggregation strategies based on the model architecture, such as channel-wise aggregation [25], [26] and layer-wise aggregation [27]. Some works mix elements from both of these solutions to achieve their goals [28], [29]. However, all existing heterogeneous federated learning methods cannot be directly used in FedRecs due to the inherent disparities between general federated learning and FedRecs as displayed in Fig 3.

- **Difference in Data Structures.** In FedRecs, a data sample is formed as $(user, item, rating)$ which is associated with specific users. Consequently, constructing a public reference dataset to facilitate knowledge distillation while ensuring privacy protection becomes impractical.
- **Difference in Model Architecture.** In FedRecs, the item embedding table plays a dominant role in determining the size of the recommendation model. The distinction between small and large models primarily lies in the dimensions of the item embedding. However, the model

⁶In this paper, heterogeneous models refer to models with different sizes.

size in federated learning is mainly related to the depth of layers and the widths of channels.

In this paper, we propose a novel federated recommendation framework named HeteFedRec, which aims to facilitate collaborative learning across clients with heterogeneous models. Since item embeddings significantly influence the entire recommendation model and user embeddings are typically private parameters not shared among clients, HeteFedRec primarily focuses on heterogeneous item embeddings. The conventional idea is to employ the padding method to aggregate item embeddings of different sizes. However, this aggregation strategy suffers from a severe mismatch problem because parameters at the same position but in different embedding tables may represent different latent factors and meanings. To address this issue, we propose unified dual-task learning, ensuring that submatrices of large embedding tables and small embedding tables share the same objective/task. Furthermore, we design a novel decorrelation regularization to prevent large item embeddings from degrading to smaller ones. Lastly, a relation-based ensemble distillation is proposed to effectively merge knowledge from embeddings of different sizes based on the intuitive idea that the ensembled spatial information derived from various types of item embeddings contains more reliable and useful knowledge. To showcase the effectiveness, efficiency, and generalization of HeteFedRec, we conduct extensive experiments on three recommendation datasets using two commonly used base recommendation models.

The main contributions of this paper are as follows:

- To the best of our knowledge, we are the first to explore model heterogeneity in federated recommender systems.
- We propose a heterogeneous model aggregation strategy tailored for FedRecs, which leverages a unified dual-task learning mechanism with dimensional decorrelation regularization to facilitate knowledge sharing among recommender models of different sizes. Furthermore, a

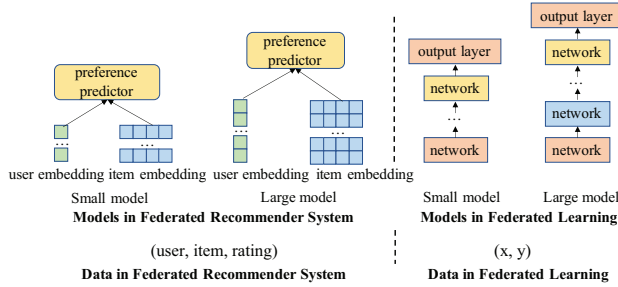


Fig. 3: The difference between federated recommender systems and federated learning in model architecture and data structure.

relation-based distillation method is proposed to distill knowledge from heterogeneous item embedding tables.

- Extensive experiments on three real-world recommendation datasets demonstrate that our proposed methods are effective and efficient and outperform all baselines.

II. RELATED WORK

In this section, we briefly review the literature on most related topics. Other involved topics such as general recommender systems [30] and federated learning can be referred to corresponding surveys [31]–[35].

A. Federated Recommender Systems

Due to the privacy-preserving ability, FedRecs have received remarkable attention recently and have been widely deployed in many recommendation scenarios [16] such as news [36], social [37], and POI recommendation [38]. The pioneering work by [12] introduced the first federated recommendation framework that utilizes federated learning with collaborative filtering models. Building upon this basic framework, extensive research has been proposed to improve the performance of FedRecs in a short time [39]. Some works aim to bridge the performance gap between FedRecs and centralized recommender systems by using advanced neural networks [14] and training techniques [40]. Others focus on reducing the training costs associated with FedRecs. For example, [13], [41] studies achieving convergence within fewer training epochs, and [42] attempts to reduce communication costs each round. Besides, some researchers propose personalized federated recommender systems [43], that attempt to balance the common-sense knowledge from global models and the personalized information from local updating. These personalized federated recommender systems aim to address the content heterogeneity in clients' data, which is orthogonal to this work.

With the notable advancements of FedRecs, the security concerns surrounding these systems have recently gained attention. Studies such as [44]–[47] demonstrate that FedRecs are susceptible to manipulation by malicious users who upload poisoned model updates, resulting in unfair and inappropriate recommendations. [48]–[50] reveal that FedRecs still suffer

potential privacy issues since adversaries can steal attribute information [51] and the system cannot forget quit clients.

Although significant progress has been made in FedRecs, existing approaches necessitate sharing a global model of the same size among clients, leading to suboptimal performance. To address this concern, this paper takes a pioneering step with HeteFedRec, enabling clients to collaboratively learn global models with varying sizes in FedRecs.

B. Heterogeneous Federated Learning

Generally, the heterogeneity research [52] in federated learning can be classified into three scenarios: data heterogeneity, device heterogeneity, and model heterogeneity. Data heterogeneity research focuses on learning a global model with clients whose training data is non-IID [20], [53], [54]. These studies attempt to strike a balance between acquiring common-sense knowledge and maintaining local personalization. For instance, Zhao et al. [55] addressed this problem by adding a small subset of global data shared among all the clients. Li et al. [56] proposed FedProx, a re-parametrization of FedAvg [10] that includes a regularization term to preserve the local model's personalization. Hermes [57] performs partial parameter averaging to retain model personalization. In addition, [58] incorporates contrastive learning to rectify model representations among individual parties. On the other hand, research on device heterogeneity considers variations in computing power among different participants [59], [60].

Our work is more related to model heterogeneity federated learning. Traditional federated learning makes an essential assumption that all clients train a model with the same architecture [35]. In model heterogeneity, this assumption is relaxed and the main challenge lies in aggregating knowledge from diverse models. One research line is to design specific aggregation strategies based on target model architecture. For example, [25], [61] design width-level aggregation for parameters with different channel scales. [27]–[29] control model size by adjusting the number of neural layers, and they only aggregate layers in the lower layers, as these layers capture fundamental features. Based on layer-level aggregation, [26], [59] applies neural pruning techniques to create heterogeneous models. Another research direction involves using knowledge distillation [24] to transfer knowledge among heterogeneous models. [21], [22] construct a public reference dataset that is accessible to all clients, leveraging soft label distillation to aggregate heterogeneous local models.

However, the aforementioned methods for model heterogeneity cannot be directly applied to FedRecs due to the notable distinctions between FedRecs and traditional federated learning. In this paper, we take the initial stride toward implementing model heterogeneity in the context of FedRecs.

C. Knowledge Distillation

Based on the types of knowledge being transferred, distillation can be classified into response-based knowledge distillation, feature-based knowledge distillation, and relation-based knowledge distillation. Response-based knowledge dis-

tillation [62], [63] transfers knowledge from a teacher model to a student model by utilizing the outputs of the final layer. Feature-based distillation, on the other hand, focuses on training a student model using features extracted from intermediate layers, as these layers learn to discriminate specific features [64]. Meanwhile, relation-based distillation involves capturing the relationships between feature maps and leveraging this knowledge to train student models [65]–[67]. In this paper, we propose a novel relation-based ensemble self-distillation method that allows for the fusion of knowledge among item embeddings of varying sizes.

III. PRELIMINARIES

A. Federated Recommendation Settings

Let $\mathcal{U} = \{u_i\}_{i=1}^{|\mathcal{U}|}$ and $\mathcal{V} = \{v_j\}_{j=1}^{|\mathcal{V}|}$ denote the sets of users and items, respectively, while $|\mathcal{U}|$ and $|\mathcal{V}|$ represent their sizes. In the context of FedRecs, each user u_i is a client that possesses a local dataset \mathcal{D}_i with the size $|\mathcal{D}_i|$ varying among users as shown in Fig. 1. Note that to protect user privacy, the private dataset \mathcal{D}_i will not be shared with any other participants. \mathcal{D}_i consists of user-item interactions (u_i, v_j, r_{ij}) , where r_{ij} is implicit feedback in this paper. $r_{ij} = 1$ indicates u_i has interacted with item v_j , and $r_{ij} = 0$ means no interaction between u_i and v_j , i.e., v_j is a negative sample. The goal of FedRecs is to predict \hat{r}_{i*} between u_i and non-interacted items and subsequently recommend the top-K ones with the highest predicted scores. To achieve this, FedRecs are trained to optimize the following formula:

$$\underset{\{\mathbf{u}_1, \dots, \mathbf{u}_i, \mathbf{V}, \Theta\}}{\text{argmin}} \sum_{u_i \in \mathcal{U}} \mathcal{L}(\mathcal{F}(\mathbf{u}_i, \mathbf{V}, \Theta) | \mathcal{D}_i) \quad (1)$$

where \mathbf{u}_i is a user embedding vector, \mathbf{V} is a $|\mathcal{V}| \times N$ item embedding matrix which is the most memory-intensive component in FedRecs. N is the embedding size (i.e., the number of dimensions). The user embedding \mathbf{u}_i is an $1 \times N$ vector. Θ is other trainable parameters such as the parameters of feedforward layers. To optimize E.q. 1 with user privacy protection, FedRecs set \mathbf{u}_i as private parameters which will only store in corresponding local devices and will not be accessed by any other participants. In contrast, \mathbf{V} and Θ are used as public parameters to achieve collaborative learning. \mathcal{F} is a base recommendation algorithm, in this paper, we tried NCF [68] and LightGCN [69] which will be introduced in the following subsection. \mathcal{L} is a loss function, in this paper, following [8], [46], [47], [49], [50], our loss function is:

$$\mathcal{L}(\mathbf{u}_i, \mathbf{V}, \Theta) = \sum_{(u_i, v_j, r_{ij}) \in \mathcal{D}_i} -r_{ij} \log \hat{r}_{ij} + (r_{ij} - 1) \log(1 - \hat{r}_{ij}) \quad (2)$$

In FedRecs, a central server coordinates all clients to optimize E.q. 1 by transmitting and aggregating the public parameters with the following federated learning protocol. In the initial stage, the central server initializes all public parameters \mathbf{V}^0 and Θ^0 , and each client independently initializes its private parameters \mathbf{u}_i^0 . After that, the following steps are iteratively executed until model convergence. At time step t , the central server first selects a group of clients

\mathcal{U}^{t-1} and distributes public parameters \mathbf{V}^{t-1} and Θ^{t-1} to these clients. The selected clients integrate the received public parameters with their private parameters \mathbf{u}_i^{t-1} , forming local recommenders, which are then trained to optimize E.q. 2 on the respective local datasets. After several epochs of local training, the client u_i sends the updates of public parameters $\nabla \mathbf{V}_i^{t-1}$ and $\nabla \Theta_i^{t-1}$ to the central server and concurrently updates its private parameters locally.

$$\mathbf{u}_i^t = \mathbf{u}_i^{t-1} - lr \nabla \mathbf{u}_i^{t-1} \quad (3)$$

where lr is the learning rate. The server employs specific strategies [10] to aggregate the updates of public parameters and use them to achieve collaborative learning:

$$\begin{aligned} \mathbf{V}^t &= \mathbf{V}^{t-1} - lr \sum_{u_i \in \mathcal{U}^{t-1}} \nabla \mathbf{V}_i^{t-1} \\ \Theta^t &= \Theta^{t-1} - lr \sum_{u_i \in \mathcal{U}^{t-1}} \nabla \Theta_i^{t-1} \end{aligned} \quad (4)$$

The above most widely used FedRec framework relies on an important assumption that all clients possess ample resources, such as sufficient training data, to support the training of recommendation models of the same sizes. However, as depicted in Fig. 1, users' data size suffers substantial variance. Consequently, the basic FedRec framework can only train a recommender system that attains suboptimal performance.

B. Base Recommendation Models

The basic FedRec framework is compatible with the majority of deep learning-based recommendation models. Among these models, NCF [68] and LightGCN [69] are two most widely used techniques. Hence, to demonstrate the generalization of our proposed methods, we utilize these two recommendation models as base models in this paper.

Neural collaborative filtering [68] (NCF) extends the method of collaborative filtering by utilizing several layers of feedforward network (FNN) to learn the complex patterns from user-item interaction data:

$$\hat{r}_{ij} = \sigma(FFN([\mathbf{u}_i, \mathbf{v}_j])) \quad (5)$$

where $FFN(\cdot)$ is a feedforward network, and its parameters are included in Θ , $[\cdot]$ denotes concatenation operation.

In LightGCN [69], users and items are treated as distinct nodes and a bipartite graph is constructed based on user-item interactions. Subsequently, a LightGCN propagation is applied on the graph to compute user and item embeddings. To ensure privacy, the propagation is only used in user's local graph. Finally, these user and item embeddings are used to predict users' preference scores via E.q. 5.

IV. METHODOLOGY: HETEFEDREC

As mentioned in Section III, the commonly used basic FedRec framework necessitates clients to share public parameters of the same sizes, a.k.a. homogeneous FedRecs. However, real-world scenarios often involve clients with varying data sizes. For clients with limited data, their model updates may be unreliable and hinder the convergence of the global

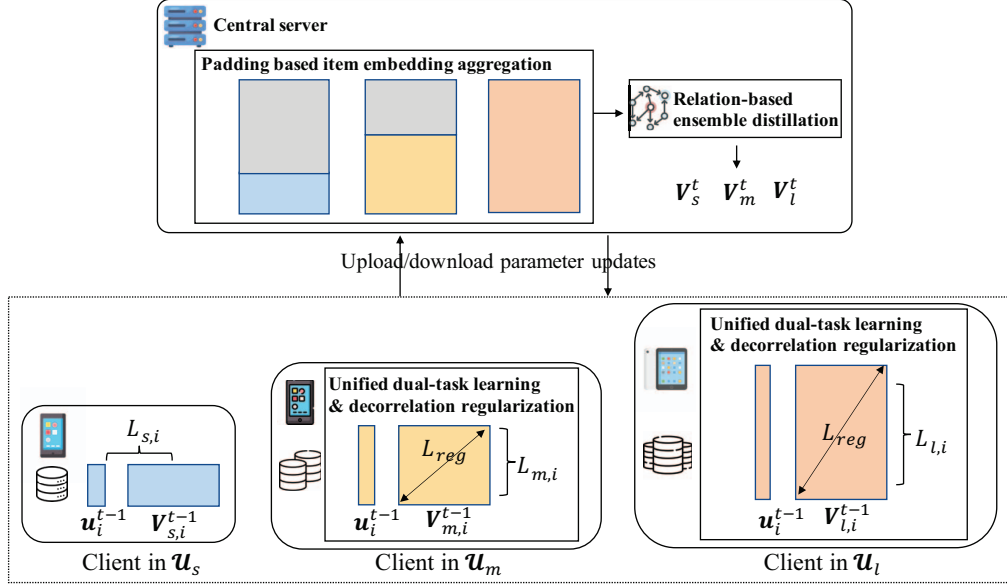


Fig. 4: The framework of HeteFedRec. To facilitate presentation, we omit Θ in this figure.

recommender system. Conversely, clients with abundant data would prefer to train larger models to obtain more accurate recommendations, benefiting from the stronger representation capabilities of larger models.

Based on the homogeneous FedRec framework, there are two naive solutions to address the aforementioned dilemma. One is to exclude the updates of public parameters from clients with limited training data and only aggregate updates from clients with sufficient data. However, as illustrated in Fig. 1, this approach would result in the exclusion of a significant portion of clients, severely compromising the recommendation performance. Another intuitive method is to opt for smaller models that are suitable for the majority of clients. Nevertheless, these smaller models have limited representation capabilities, thereby leading to unsatisfactory recommendation performance.

In this paper, we introduce HeteFedRec, the first heterogeneous federated recommender system. HeteFedRec allows for the customization of the recommendation model size for each client based on their available resources. Clients with limited data can train smaller models, while clients with abundant data can still benefit from the services provided by larger models. In the following subsections, we will first introduce the fundamental settings of our HeteFedRec, which builds upon the basic FedRec framework. Subsequently, we will present the technical details of HeteFedRec, including (1) a heterogeneous recommendation model aggregation strategy that contains a unified dual-task learning and a dimensional decorrelation regularization and (2) a relation-based ensemble self-distillation. Fig. 4 presents an overview of HeteFedRec.

A. Settings of HeteFedRec

In HeteFedRec, the size of a client's recommendation model can be customized based on its specific situation. In this paper, following [28], without loss of generality, we categorize

clients into three groups based on the scale of user-item interactions: small clients \mathcal{U}_s , medium clients \mathcal{U}_m , and large clients \mathcal{U}_l , to convenient the introduction of our proposed methods. Each group is assigned corresponding public parameters $\{\mathbf{V}, \Theta\}_{s,m,l}$. Here, $\mathbf{V}_s, \mathbf{V}_m, \mathbf{V}_l$ are item embeddings of sizes $|\mathcal{V}| \times N_s, |\mathcal{V}| \times N_m, |\mathcal{V}| \times N_l$, respectively, with the relationship $N_s < N_m < N_l$. The size of Θ depends on the dimension length of \mathbf{V} as it primarily includes parameters in feedforward layers. Note that the settings of division and dimension sizes are hyper-parameters and can be adjusted according to specific situations.

In general, item embedding table \mathbf{V} dominates the number of parameters in the entire recommender system. With these settings, users in \mathcal{U}_s will have smaller models, while users in \mathcal{U}_l will have larger models. Under this heterogeneous circumstance, FedRecs are trained to optimize the following objective:

$$\underset{\{\mathbf{u}_1 \dots \mathbf{u}_l, \{\mathbf{V}, \Theta\}_{s,m,l}\}}{\operatorname{argmin}} \sum_{a \in \{s,m,l\}} \sum_{u_i \in \mathcal{U}_a} \mathcal{L}(\mathcal{F}(\mathbf{u}_i \mathbf{V}_a \Theta_a) | \mathcal{D}_i) \quad (6)$$

One straightforward approach to optimize E.q. 6 is using a clustered-aggregation method. In this approach, the central server aggregates clients' model updates separately based on their respective model sizes. However, this method falls short of achieving optimal performance due to two main reasons: (1) the knowledge cannot be effectively shared across different groups, hindering the ability to leverage collaborative information; (2) since each training sample can only update the corresponding item embedding vector, the item embedding table becomes data-hungry, and it may not be well-trained within each isolated group, leading to suboptimal performance.

Another method is to aggregate these heterogeneous models based on some observations, which has proven effective in general federated learning. For example, HeteroFL [25] di-

rectly conducts parameter-averaging over heterogeneous models based on the observation that deep and wide neural networks can drop a tremendous number of parameters per layer. [27], [28] utilize layer-wise aggregation, focusing on the lower layers that retain more pre-trained and fundamental knowledge. However, all these preconditions are infeasible in the context of FedRecs, rendering these methods inapplicable. In light of this, HeteFedRec introduces a novel heterogeneous aggregation tailored for federated recommender systems.

B. Heterogeneous Recommendation Model Aggregation

In FedRecs, the private parameters, specifically user embeddings, are updated locally and do not communicate with other clients. As a result, the update of these private parameters can still adhere to the protocol outlined in the basic FedRec framework. Consequently, HeteFedRec primarily focuses on the knowledge aggregation of the public parameters, namely $\{\mathbf{V}, \Theta\}_{s,m,l}$. Among these parameters, the item embeddings $\{\mathbf{V}\}_{s,m,l}$ play a particularly crucial role as they dominate the number of parameters in a recommender system. In what follows, we will first introduce the aggregation strategy design for item embeddings and then provide a brief overview of the aggregation of Θ .

Padding based Heterogeneous Item Embedding Aggregation. At epoch t , for the uploaded item embedding updates, HeteFedRec's central server first pads the smaller item embedding updates to align them with the largest one using the following function:

$$\nabla_p \mathbf{V} \leftarrow padding(\nabla \mathbf{V} | N_l) \quad (7)$$

where $padding(\cdot | N_l)$ is a padding function that transforms the matrix updates $\nabla \mathbf{V}$ to $\nabla_p \mathbf{V}$ of size $|\mathcal{V}| \times N_l$ by filling $\mathbf{0}$ vectors. After padding, the central server performs the summation operation on all received item embedding updates and the aggregated updates for each type of embedding are as follows:

$$\begin{aligned} \nabla \mathbf{V}_{agg}^{t-1} &= \sum_{u_i \in \mathcal{U}_s^{t-1}} \nabla_p \mathbf{V}_{s,i}^{t-1} + \sum_{u_i \in \mathcal{U}_m^{t-1}} \nabla_p \mathbf{V}_{m,i}^{t-1} + \sum_{u_i \in \mathcal{U}_l^{t-1}} \nabla \mathbf{V}_{l,i}^{t-1} \\ \nabla \mathbf{V}_s^{t-1} &= \nabla \mathbf{V}_{agg}^{t-1} \\ \nabla \mathbf{V}_m^{t-1} &= \nabla \mathbf{V}_{agg}^{t-1} \\ \nabla \mathbf{V}_l^{t-1} &= \nabla \mathbf{V}_{agg}^{t-1} \end{aligned} \quad (8)$$

where \mathcal{U}_*^{t-1} is the selected clients in the group \mathcal{U}_* at epoch t . $\nabla \mathbf{V}_{*}^{t-1}$ refers to a submatrix of size $|\mathcal{V}| \times N_x$.

Then, the aggregated updates are used to modify corresponding item embeddings:

$$\begin{aligned} \mathbf{V}_s^t &= \mathbf{V}_s^{t-1} - lr \nabla \mathbf{V}_s^{t-1} \\ \mathbf{V}_m^t &= \mathbf{V}_m^{t-1} - lr \nabla \mathbf{V}_m^{t-1} \\ \mathbf{V}_l^t &= \mathbf{V}_l^{t-1} - lr \nabla \mathbf{V}_l^{t-1} \end{aligned} \quad (9)$$

Note that we set the same learning rate lr for all types of item embeddings since we do not want to introduce too many hyperparameters and we find this setting is effective enough.

According to the aggregation, when HeteFedRec initializes \mathbf{V}_s^0 , \mathbf{V}_m^0 and \mathbf{V}_l^0 from the same point respectively, we can get following relationship:

$$\begin{aligned} \mathbf{V}_s^t &= \mathbf{V}_{m[N_s]}^t = \mathbf{V}_{l[N_s]}^t \\ \mathbf{V}_m^t &= \mathbf{V}_{l[N_m]}^t \end{aligned} \quad (10)$$

Unified Dual-task Learning. However, the above naive aggregation cannot effectively train usable FedRecs because it suffers from severe mismatch problems of updates. Specifically, during clients' local training, the updates of $\nabla \mathbf{V}_{s,i}^{t-1}$, $\nabla \mathbf{V}_{m,i}^{t-1}$, and $\nabla \mathbf{V}_{l,i}^{t-1}$ are calculated by optimizing the loss function E.q. 2. Hence, as a part of these updates, the meaning of the subparts of updates, e.g., $\nabla \mathbf{V}_{m[N_s],i}^{t-1}$, $\nabla \mathbf{V}_{l[N_s],i}^{t-1}$, and $\nabla \mathbf{V}_{l[N_m],i}^{t-1}$, is unclear. In other words, the subparts of updates, $\nabla \mathbf{V}_{m[N_s],i}^{t-1}$ and $\nabla \mathbf{V}_{l[N_s],i}^{t-1}$, or $\nabla \mathbf{V}_{l[N_m],i}^{t-1}$, are not computed to guide $\mathbf{V}_{s,i}^{t-1}$ or $\mathbf{V}_{m,i}^{t-1}$ to solve the recommendation problem, therefore, it would be difficult to get convergence by aggregating using E.q. 8.

Inspired by [29] that leverages multiple the same objective functions to optimize a deep model's submodel with different depths, HeteFedRec employs a unified dual-task learning mechanism to ensure submatrices within larger item embeddings share the same objective as small item embeddings. Specifically, each user u_i will utilize one of the following loss functions depending on the group $\mathcal{U}_{s,m,l}$ it belongs to:

$$\begin{aligned} \mathcal{L}_{s,i} &= \mathcal{L}(\mathbf{u}_i^{t-1}, \mathbf{V}_{s,i}^{t-1}, \Theta_{s,i}^{t-1}) \\ \mathcal{L}_{m,i} &= \mathcal{L}(\mathbf{u}_{i[N_s]}^{t-1}, \mathbf{V}_{m[N_s],i}^{t-1}, \Theta_{s,i}^{t-1}) + \mathcal{L}(\mathbf{u}_i^{t-1}, \mathbf{V}_{m,i}^{t-1}, \Theta_{m,i}^{t-1}) \\ \mathcal{L}_{l,i} &= \mathcal{L}(\mathbf{u}_{i[N_s]}^{t-1}, \mathbf{V}_{l[N_s],i}^{t-1}, \Theta_{s,i}^{t-1}) + \mathcal{L}(\mathbf{u}_{i[N_m]}^{t-1}, \mathbf{V}_{l[N_m],i}^{t-1}, \Theta_{m,i}^{t-1}) \\ &\quad + \mathcal{L}(\mathbf{u}_i^{t-1}, \mathbf{V}_{l,i}^{t-1}, \Theta_{l,i}^{t-1}) \end{aligned} \quad (11)$$

where $\mathcal{L}(\mathbf{u}_i, \mathbf{V}, \Theta)$ is the loss function defined in E.q. 2. Note that in order to calculate E.q. 11, clients in \mathcal{U}_m will not only receive Θ_m but also Θ_s , meanwhile, users in \mathcal{U}_l will have all the Θ_s , Θ_m and Θ_l . Since Θ mainly comprises parameters from feedforward networks with 2 or 3 layers, these extra communication costs will be negligible. Based on the above equation, larger item embeddings' submatrices updates, such as $\nabla \mathbf{V}_{m[N_s],i}^{t-1}$, $\nabla \mathbf{V}_{l[N_s],i}^{t-1}$, and $\nabla \mathbf{V}_{l[N_m],i}^{t-1}$, can be directly aggregated to optimize smaller item embeddings since these updates are trained to minimize the same recommendation loss function. Fig. 5 illustrates the unified dual-task learning.

Dimensional Decorrelation Regularization. Although E.q. 11 fixes the updates mismatch problem, it may introduce a new challenge known as dimensional collapse for larger embeddings \mathbf{V}_m and \mathbf{V}_l . Specifically, by solely optimizing the submatrices $\mathbf{V}_{l[N_s],i}$ and $\mathbf{V}_{m[N_s],i}$, all terms in \mathcal{L}_l or \mathcal{L}_m are reduced. This phenomenon causes the item representation of larger models to be primarily confined to the low-dimensional space, rendering the remaining parameters $\mathbf{V}_{m[N_s],i}$ and $\mathbf{V}_{l[N_s],i}$ redundant. Consequently, HeteFedRec degrades to homogeneous FedRecs that solely rely on \mathbf{V}_s , thus diminishing the benefits of model heterogeneity.

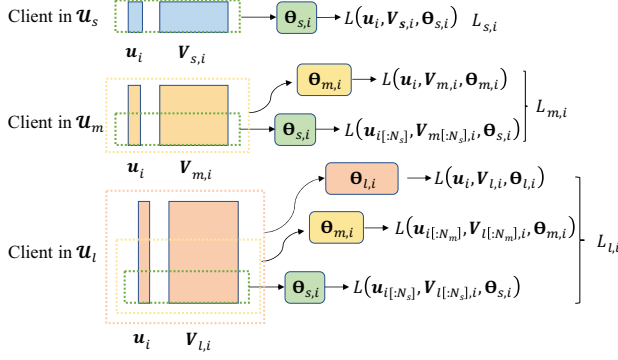


Fig. 5: Overview of unified dual-task learning.

To maintain the importance of each dimension in a matrix, one natural way is to retain their singular values of the covariance matrix. Therefore, to prevent $\mathbf{V}_{m,i}$ and $\mathbf{V}_{l,i}$ collapsing to small embedding $\mathbf{V}_{s,i}^t$, we can add the regularization term $\mathcal{L}_{singular}(\mathbf{V}_{m,i}^{t-1})$ and $\mathcal{L}_{singular}(\mathbf{V}_{l,i}^{t-1})$ in E.q. 11 to penalize the variance among each dimension's singular value:

$$\mathcal{L}_{singular}(\mathbf{V}) = \frac{1}{N} \sum_{i=1}^N (\lambda_i - \frac{1}{N} \sum_{j=1}^N \lambda_j)^2 \quad (12)$$

where λ is the singular value of \mathbf{V} 's covariance matrix. However, the computation of singular values is expensive. [70], [71] have demonstrated that instead of penalizing the variance of singular values, regularizing the Frobenius norm of the correlation matrix can achieve the same effects. Thus, in this paper, we employ the following decorrelation regularization term to ensure that $\mathbf{V}_{l[N_s],i}^t$ or $\mathbf{V}_{m[N_s],i}^t$ encode unique knowledge:

$$\mathcal{L}_{reg}(\mathbf{V}) = \frac{1}{N} \left\| \text{corr}\left(\frac{\mathbf{V} - \bar{\mathbf{V}}}{\sqrt{\text{var}(\mathbf{V})}}\right) \right\|_F \quad (13)$$

$\|\cdot\|_F$ is Frobenius norm. $\text{corr}(\cdot)$ calculates the correlation matrix. $\bar{\mathbf{V}}$ is the columns' mean of \mathbf{V} by default. $\text{var}(\cdot)$ computes the variance of a matrix.

Then, the loss function for clients in groups \mathcal{U}_m and \mathcal{U}_l is transformed from E.q. 11 to:

$$\begin{aligned} \mathcal{L}'_{m,i} &= \mathcal{L}_{m,i} + \alpha \mathcal{L}_{reg}(\mathbf{V}_{m,i}^{t-1}) \\ \mathcal{L}'_{l,i} &= \mathcal{L}_{l,i} + \alpha \mathcal{L}_{reg}(\mathbf{V}_{l,i}^{t-1}) \end{aligned} \quad (14)$$

where α is the factor that controls the importance of the regularization term. We adopt the same α for both medium and large item embeddings to minimize the number of hyperparameters. Through experiments, we have observed that this setting proves to be sufficiently effective.

Now, after clients finish local training with the loss function E.q. 14, HeteFedRec's central server can conduct heterogeneous aggregation for item embeddings among different models using E.q. 8.

$\{\Theta\}_{s,m,l}$ **Aggregation.** For Θ , we simply aggregate them with the same size:

$$\begin{aligned} \Theta_s^t &= \Theta_s^{t-1} - lr \sum_{u_i \in \mathcal{U}_s^{t-1} \cup \mathcal{U}_m^{t-1} \cup \mathcal{U}_l^{t-1}} \nabla \Theta_{s,i}^{t-1} \\ \Theta_m^t &= \Theta_m^{t-1} - lr \sum_{u_i \in \mathcal{U}_m^{t-1} \cup \mathcal{U}_l^{t-1}} \nabla \Theta_{m,i}^{t-1} \\ \Theta_l^t &= \Theta_l^{t-1} - lr \sum_{u_i \in \mathcal{U}_l^{t-1}} \nabla \Theta_{l,i}^{t-1} \end{aligned} \quad (15)$$

Algorithm 1 HeteFedRec: Federated Recommender System with Model Heterogeneity.

Input: global epoch T ; local epoch L ; learning rate lr, \dots

Output: public parameters $\{\mathbf{V}, \Theta\}_{s,m,l}$, local client embeddings $\mathbf{u}_i |_{i \in \mathcal{U}}$

- 1: Initialize public parameters $\{\mathbf{V}^0, \Theta^0\}_{s,m,l}$
- 2: **for** each round $t=1, \dots, T$ **do**
- 3: sample a fraction of clients \mathcal{U}^{t-1} from \mathcal{U}
- 4: **for** $u_i \in \mathcal{U}^{t-1}$ **in parallel do**
- 5: // execute on client sides
- 6: CLIENTTRAIN(u_i)
- 7: **end for**
- 8: // the server executes heterogeneous aggregation
- 9: $\{\nabla \mathbf{V}^{t-1}\}_{s,m,l} \leftarrow$ aggregate updates using E.q. 8
- 10: $\{\mathbf{V}^t, \Theta^t\}_{s,m,l} \leftarrow$ update public parameters using E.q. 9 and E.q. 15
- 11: // the server executes distillation
- 12: $\{\mathbf{V}^t\}_{s,m,l} \leftarrow$ update embeddings using E.q. 17
- 13: **end for**
- 14: **function** CLIENTTRAIN(u_i)
- 15: download public parameters from the server
- 16: **if** $u_i \in \mathcal{U}_s$ **then**
- 17: calculate $\nabla \mathbf{u}_i^{t-1}, \nabla \mathbf{V}_{s,i}^{t-1}, \nabla \Theta_{s,i}^{t-1}$ using $\mathcal{L}_{s,i}$ in E.q. 11
- 18: upload $\nabla \mathbf{V}_{s,i}^{t-1}, \nabla \Theta_{s,i}^{t-1}$ to the server
- 19: **else if** $u_i \in \mathcal{U}_m$ **then**
- 20: calculate $\nabla \mathbf{u}_i^{t-1}, \nabla \mathbf{V}_{m,i}^{t-1}, \{\nabla \Theta_i^{t-1}\}_{s,m}$ using $\mathcal{L}'_{m,i}$ in E.q. 14
- 21: upload $\nabla \mathbf{V}_{m,i}^{t-1}, \{\nabla \Theta_i^{t-1}\}_{s,m}$ to the server
- 22: **else** // $u_i \in \mathcal{U}_l$
- 23: calculate $\nabla \mathbf{u}_i^{t-1}, \nabla \mathbf{V}_{l,i}^{t-1}, \{\nabla \Theta_i^{t-1}\}_{s,m,l}$ using $\mathcal{L}'_{l,i}$ in E.q. 14
- 24: upload $\nabla \mathbf{V}_{l,i}^{t-1}, \{\nabla \Theta_i^{t-1}\}_{s,m,l}$ to the server
- 25: **end if**
- 26: $\mathbf{u}_i^t \leftarrow$ update local private parameter using E.q. 3
- 27: **end function**

C. Relation-based Ensemble Self Knowledge Distillation

Knowledge distillation has been used to facilitate knowledge sharing among heterogeneous models in federated learning [52]. However, these works [21], [22], [28] require the construction of publicly available reference datasets. In the context of FedRecs, where data samples are sensitive and

specific to users, these knowledge distillation methods cannot be directly applied to heterogeneous FedRecs.

In this paper, we propose a novel relation-based ensemble self-distillation mechanism to further enhance the performance of heterogeneous FedRecs, eliminating the need for constructing a reference dataset. The basic idea of our method is that, when item embeddings are well-trained from collaborative information, the similarity of items should be consistent among different embedding sizes. Therefore, we transfer knowledge by harmonizing the spatial information across item embeddings of different sizes, i.e., keep the relative distance relationship of item vectors consistent among item embedding tables.

Specifically, after getting $\{\mathbf{V}^t\}_{s,m,l}$ using heterogeneous recommendation model aggregation at epoch t , the central server conducts distillation with the following steps. First, to avoid heavy computation costs, it randomly selects a subset of items \mathcal{V}_{kd} from the whole item set \mathcal{V} as the target distillation items. Then, it calculates selected items' distance from each other using $\{\mathbf{V}^t\}_{s,m,l}$ and obtains the ensemble distance by averaging:

$$d_{ens}(\mathcal{V}_{kd}) = \frac{1}{3} \sum_{a \in \{s,m,l\}} d(\mathbf{V}_a^t, \mathcal{V}_{kd}) \quad (16)$$

where $d(\cdot)$ calculates the distance among selected items. In this paper, we use cosine similarity as the distance function. Finally, for each embedding table in $\{\mathbf{V}^t\}_{s,m,l}$, the central server computes distillation loss as follows:

$$\mathcal{L}_{kd}(\mathbf{V}) = \|d(\mathbf{V}, \mathcal{V}_{kd}) - d_{ens}(\mathcal{V}_{kd})\|_2^2 \quad (17)$$

Algorithm 1 summarizes HeteFedRec using pseudocode.

V. EXPERIMENTS

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

- **RQ1.** How effective is our HeteFedRec compared to homogeneous and heterogeneous FedRec baselines?
- **RQ2.** How efficient is our HeteFedRec compared to the baselines in terms of model convergence speed and communication cost?
- **RQ3.** How does HeteFedRec benefit from each key component?
- **RQ4.** How does the ratio of client division affect HeteFedRec's performance?
- **RQ5.** How does the model size affect HeteFedRec's performance?
- **RQ6.** How does the value of hyperparameter α affect HeteFedRec's performance?

A. Datasets

We conduct extensive experiments on three real-world recommendation datasets: MovieLens-1M (ML) [17], Anime, and Douban [72], covering three different scenarios (movie recommendation, anime recommendation, and book recommendation) to evaluate the performance of our HeteFedRec. The statistics of these datasets are presented in Table I. "Avg."

is the average number of user-item interactions. "< 50%" or "< 80%" represent 50% or 80% of users' interaction numbers that are less than certain values. In this paper, we use these values to divide clients into $\mathcal{U}_{s,m,l}$ without specific notation. The ML dataset consists of users' ratings for movies and includes 6,040 users, 3,706 items, and 1,000,209 interactions. The Anime dataset contains user preference data crawled from MyAnimeList⁷ and consists of 10,482 users, 6,888 items, and 1,265,530 watching records. The Douban dataset is a subset from Douban book, comprising 330,268 interactions between 1,833 users and 7,397 items. Following the settings of recommendation with implicit feedback [44], [49], [68], we binarize all datasets' user ratings, transforming all ratings to $r_{ij} = 1$, and negative instances are sampled with a ratio of 1 : 4. For each dataset, 80% of data and 20% of data are used as training and test set. When a client is selected for training, 10% of its training data will be used as the validation set to guide the local training. Note that to protect data privacy, following [44], [49], [68], a user is a client, and users can only use their own data during the training process.

TABLE I: Statistics of recommendation datasets

Dataset	Users	Items	Interactions	Avg.	< 50%	< 80%
ML	6,040	3,706	1,000,209	165	77	203
Anime	10,482	6,888	1,265,530	120	69	150
Douban	1,833	7,397	330,268	180	115	244

B. Evaluation Metrics

To evaluate the effectiveness of the proposed heterogeneous federated recommendation framework, we employ the widely adopted metrics [69], [73] Recall at Rank 20 (Recall@20) and Normalized Discounted Cumulative Gain at Rank 20 (NDCG@20) to evaluate the recommender system's performance. Recall measures the average probability of relevant items being successfully recommended to users, while NDCG takes into account the position of relevant items in the recommendation list. These metrics provide a comprehensive evaluation of the recommender system's performance.

C. Baselines

Since this is the first work on federated recommender systems with heterogeneous model sizes, we construct the following baselines for comparison:

- **All Small.** This baseline deploys small recommendation models to all clients without considering their data amounts. While this approach avoids the dilemma of weak clients not being able to support larger models, it limits the overall ability of FedRecs.
- **All Large.** This baseline deploys large recommendation models to all clients without considering their data amounts. However, since clients with small amounts of data cannot effectively support the training of large models, their updates can even negatively impact the overall performance of FedRecs.

⁷<https://myanimelist.net/>

- **All Large/Exclusive.** This baseline deploys large recommendation models to all clients but only aggregates updates from clients with sufficient data.
- **Standalone.** This baseline assigns heterogeneous recommendation models to clients based on their data amounts. Each client independently trains its model without collaboration or knowledge sharing with other clients.
- **Clustered FedRec.** Following the clustered scheme in federated learning [74], [75], we construct Clustered FedRec. It aggregates the same type of model.
- **Directly Aggregate.** This method directly aggregates heterogeneous recommendation model updates using E.q. 8. However, since these updates are not guided by the unified dual-task learning mechanism, it results in a significant updates mismatch problem.

In general, the first three baselines, “All Small”, “All Large”, and “All Large/Exclusive”, are homogeneous FedRecs, where the recommendation models are either all small or all large, regardless of the client’s data amounts. On the other hand, the last three baselines, “Standalone”, “Clustered FedRec”, and “Directly Aggregate”, are heterogeneous FedRecs inspired by corresponding federated learning schemes. These heterogeneous FedRecs consider the varying data amounts of clients and aim to address the challenges posed by resource heterogeneity in the federated recommender system setting.

D. Parameter Settings

On ML and Anime datasets, the dimensions N_s , N_m , and N_l are 8, 16, 32 for both Fed-NCF and Fed-LightGCN. For the Douban, N_s , N_m , and N_l are set to 32, 64, and 128 for both Fed-NCF and Fed-LightGCN as users in Douban have a higher average number of interactions than the other datasets. The proportion of \mathcal{U}_s , \mathcal{U}_m , and \mathcal{U}_l is 5 : 3 : 2 without specific noting. Both Fed-NCF and Fed-LightGCN models utilize three feedforward layers with $[2 \times N_s, 8, 8]$ dimensions for different types of models. The layer of LightGCN propagation is 1. Adam [76] with 0.001 learning rate is adopted as the optimizer. At the beginning of an epoch, the server shuffles the queue of clients. Then, at each epoch, there are several rounds for the central server to traverse the client queue. During each round, the central server selects 256 users for training.

E. The Effectiveness of HeteFedRec (RQ1)

Table II presents the comparison of HeteFedRec with six baselines. Among homogeneous baselines, we can see that “All Small” achieves better performance than “All Large”. This phenomenon supports our argument that clients with insufficient data cannot support large models’ training therefore even with stronger models, “All Large” cannot achieve better performance. Besides, we can observe a significant performance drop of the method “All Large/Exclusive”, since it excludes the updates from \mathcal{U}_s which represent a large proportion of the entire client set. In heterogeneous FedRecs, “Standalone” exhibits the poorest performance due to the lack of collaboration among clients. The clustered scheme, which

has been used in federated learning to address model heterogeneity, proves ineffective for training a usable recommender system in “Clustered FedRec”. This could be attributed to the fact that collaborative information is crucial in recommender systems, and the clustered scheme hinders collaborative learning among clients. The “Directly Aggregate” method attempts to fuse knowledge between heterogeneous models but suffers from severe updates mismatch problems. Essentially, “Directly Aggregate” can be viewed as HeteFedRec removed unified dual-task learning, dimensional decorrelation regularization, and knowledge distillation components. In baselines, all heterogeneous FedRecs fail to train a FedRec with comparable performance to homogeneous baselines. As the first heterogeneous method tailored for FedRecs, our HeteFedRec outperforms all heterogeneous and homogeneous baselines on all datasets using both Fed-NCF and Fed-LightGCN by a significant margin, demonstrating the effectiveness of our proposed methods.

To gain further insight into the performance gap, we analyze the NDCG scores in detailed user groups as shown in Fig. 6. We focus on analyzing the performance of “All Small”, “All Large”, and HeteFedRec, as these two homogeneous baselines achieve performances closest to our HeteFedRec. Overall, the performance of all three methods follows a consistent trend: NDCG scores are higher in \mathcal{U}_l and \mathcal{U}_m compared to \mathcal{U}_s . This is because that \mathcal{U}_l and \mathcal{U}_m have more training data available compared to \mathcal{U}_s . Specifically, on ML and Anime datasets, “All Small” outperforms “All Large” in \mathcal{U}_s for both Fed-NCF and Fed-LightGCN. However, in \mathcal{U}_l , the performance of “All Small” is generally worse than “All Large”. This intriguing observation suggests that clients in \mathcal{U}_s struggle to train large models effectively, while clients in \mathcal{U}_l benefit from the larger models with stronger learning abilities, aligning with our motivation to introduce model heterogeneity in FedRecs. On the Douban dataset, “All Small” consistently outperforms “All Large” in all user groups. This may be attributed to the large model size (i.e., 128 dimension size) of the embedding tables, making it challenging for all clients to handle the training. Our HeteFedRec achieves the best performance across all client groups and scenarios. This can be attributed to two factors: (1) the assignment of models of different sizes based on clients’ data amounts, enabling clients to train appropriate models of suitable sizes, and (2) the effective knowledge transfer among heterogeneous recommendation models through our proposed heterogeneous aggregation strategies, leading to performance improvements in all user groups.

F. The Efficiency of HeteFedRec (RQ2)

Fig. 7 illustrates the convergence of HeteFedRec compared to “All Small” and “All Large” as these two baselines achieved comparable performance. Due to the limitation of space, we only present the results on the ML dataset, similar trends can be observed on other datasets. Generally, HeteFedRec achieves better performance after about 10 epochs and 20 epochs in Fed-NCF and Fed-LightGCN respectively. “All Small” is the method that reaches convergence within the fewest epochs

TABLE II: The comparison of the overall performance of HeteFedRec and baselines. The best values on each dataset are bold. Underline indicates the best values in homogeneous methods.

	Type	Methods	ML		Anime		Douban	
			Recall	NDCG	Recall	NDCG	Recall	NDCG
Fed-NCF	Homogeneous	All Small	0.02203	0.04328	0.04301	0.04962	0.00759	0.01087
		All Large	<u>0.02558</u>	0.04028	0.02727	0.04442	0.00726	0.00878
		All Large/Exclusive	0.00956	0.01753	0.01199	0.02458	0.00702	0.00856
	Heterogeneous	Standalone	0.00615	0.01108	0.00279	0.00411	0.00209	0.00295
		Clustered FedRec	0.01712	0.02235	0.01508	0.01581	0.00248	0.00501
		Directly Aggregate	0.01177	0.02207	0.01903	0.03151	0.00247	0.00502
		HeteFedRec(Ours)	0.02662	0.04781	0.05855	0.05655	0.01101	0.01290
Fed-LightGCN	Homogeneous	All Small	0.02251	0.04232	0.02924	0.04824	0.00350	0.00530
		All Large	<u>0.02301</u>	0.04197	0.02825	0.04788	0.00234	0.00378
		All Large/Exclusive	0.00924	0.01891	0.01702	0.01467	0.00215	0.00363
	Heterogeneous	Standalone	0.00605	0.01085	0.00278	0.00411	0.00190	0.00263
		Clustered FedRec	0.01483	0.02633	0.01443	0.01379	0.00259	0.00480
		Directly Aggregate	0.01454	0.02657	0.01450	0.01437	0.00257	0.00479
		HeteFedRec(Ours)	0.02434	0.04313	0.03306	0.05177	0.00393	0.00639

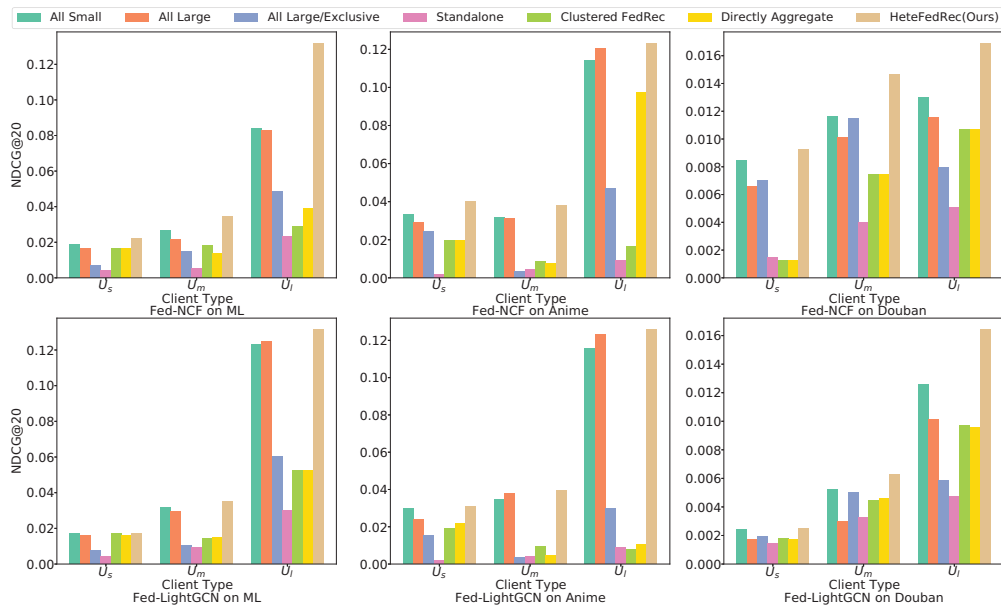


Fig. 6: Detailed performance of “All Small”, “All Large”, and “HeteFedRec(Ours)” in different user groups.

TABLE III: The cost of one-time transmitting between clients and the central server.

Client Type	All Small	All Large	HeteFedRec
\mathcal{U}_s	$\text{size}(\mathbf{V}_s + \Theta_s)$	$\text{size}(\mathbf{V}_l + \Theta_l)$	$\text{size}(\mathbf{V}_s + \Theta_s)$
\mathcal{U}_m	$\text{size}(\mathbf{V}_s + \Theta_s)$	$\text{size}(\mathbf{V}_l + \Theta_l)$	$\text{size}(\mathbf{V}_m + \{\Theta\}_{s,m})$
\mathcal{U}_l	$\text{size}(\mathbf{V}_s + \Theta_s)$	$\text{size}(\mathbf{V}_l + \Theta_l)$	$\text{size}(\mathbf{V}_l + \{\Theta\}_{s,m,l})$

because it averagely has the smallest recommendation models and small models are usually easy to get convergence. The convergence speed for “All Large” and our HeteFedRec is similar, but the converged performance of HeteFedRec is much better than these two baselines. In conclusion, HeteFedRec achieves convergence within a reasonable number of epochs compared to “All Small” and “All Large”.

The communication costs are another important consideration when using a FedRec framework. Table III compares HeteFedRec’s communication costs with “All Small” and




Fig. 7: The performance trend during the training process for Fed-NCF and Fed-LightGCN on ML. A similar trend can be seen in Anime and Douban.

“All Large” formally. The only additional costs incurred by HeteFedRec are $\text{size}(\Theta_s)$ for clients in \mathcal{U}_m and $\text{size}(\Theta_{s,m})$

TABLE IV: Ablation Study. “RESKD” is short for “Relation-based Ensemble Self Knowledge Distillation”, “DDR” is abbreviation for “Dimensional Decorrelation Regularization”, and “UDL” indicates “Unified Dual-task Learning”.

		ML		Anime		Douban	
		Recall	NDCG	Recall	NDCG	Recall	NDCG
Fed-NCF	HeteFedRec	0.02662	0.04781	0.05855	0.05655	0.01101	0.01290
	- RESKD	0.02620	0.04636	0.05696	0.05504	0.00956	0.01232
	- RESKD,DDR	0.02390	0.04332	0.05766	0.05355	0.00797	0.01167
	- RESKD,DDR,UDL	0.01177	0.02207	0.01903	0.03151	0.00247	0.00502
Fed-LightGCN	HeteFedRec	0.02434	0.04313	0.03306	0.05177	0.00393	0.00639
	- RESKD	0.02332	0.04307	0.03282	0.05072	0.00372	0.00618
	- RESKD,DDR	0.02231	0.04225	0.03129	0.04812	0.00351	0.00563
	- RESKD,DDR,UDL	0.01454	0.02657	0.01450	0.01437	0.00257	0.00479

TABLE V: The variance of singular values in the covariance matrix of the largest item embedding \mathbf{V}_l . A higher value implies a more severe dimensional collapse.

		ML	Anime	Douban
Fed-NCF	- DDR	0.4573	0.9190	0.0523
	+ DDR	0.0974	0.0838	0.0167
Fed-LightGCN	- DDR	0.0459	0.0421	0.0348
	+ DDR	0.0208	0.0240	0.0171

for users in \mathcal{U}_l . Considering our experimental settings and taking the ML dataset as an example, the parameter numbers of \mathbf{V}_s , \mathbf{V}_m , and \mathbf{V}_l are 29648, 59296 and 118592, and the parameters of $\{\Theta_{s,m,l}\}$ are generally from feedforward layers and their sizes are no more than one hundred. Therefore, these additional costs are negligible.

G. Ablation Study (RQ3)

HeteFedRec comprises three crucial components: a unified dual-task learning to address mismatch problems, a dimensional decorrelation regularization to prevent dimensional collapse, and a relation-based self-knowledge distillation to further merge knowledge among heterogeneous models. In this part, we investigate the impact of each component. Specifically, we gradually remove knowledge distillation (RESKD), regularization (DDR), and unified dual-task learning (UDL) and present the experimental results in Table IV. When RESKD is removed, there is a slight drop in the performance of HeteFedRec, indicating the effectiveness of RESKD. We then continue by removing the regularization term and find that, HeteFedRec’s performance becomes similar to “All Small” in Table II. This is because, without the dimensional decorrelation regularization term, HeteFedRec experiences a dimensional collapse problem, resulting in a degradation of its model’s representation ability to that of small models. When all three components are eliminated, HeteFedRec becomes equivalent to “Directly Aggregate”, and the recommender system’s performance is significantly decreased, highlighting the crucial role of our unified dual-task learning mechanism.

To further understand the effectiveness of DDR, we calculate the variance of singular values in the covariance matrices of the largest item embeddings \mathbf{V}_l , as shown in Table V. A smaller variance indicates a more balanced importance among dimensions, suggesting the alleviation of the dimensional

collapse problem. As displayed in Table V, after applying DDR, the variance of singular values is reduced in all cases.

H. Impact of Client Division (RQ4)

In this part, we thoroughly investigate the influence of client division. Specifically, we conduct experiments with three different division ratios 5 : 3 : 2, 1 : 1 : 1, and 2 : 3 : 5. The first division setting is the original one used in the main experiments, which assumes that most clients cannot afford the training of larger models. The 1 : 1 : 1 ratio is a neutral strategy that evenly distributes clients into the \mathcal{U}_s , \mathcal{U}_m , and \mathcal{U}_l groups based on the number of users’ training data. The 2 : 3 : 5 ratio is an optimistic division solution that assumes only 20% of clients lack training data and most clients have sufficient data to train stronger recommender systems. Table VI presents the experimental results of HeteFedRec under different client divisions. We also include “All Small” and “All Large” in Table VI for comparison, as they roughly correspond to the division ratios of 10 : 0 : 0 and 0 : 0 : 10, respectively. Overall, HeteFedRec with the conservative ratio setting (i.e., 5 : 3 : 2) achieves better performance among these three division strategies. This result is consistent with the interaction number distribution of ML, Anime, and Douban as indicated in Fig. 1, where most users have insufficient interaction data. In addition, as we move from the left to the right in Table VI (i.e., 5 : 3 : 2 to “All Large”), more and more clients are assigned with larger models, but the overall performance continues to deteriorate. This observation further supports the importance of model size heterogeneity in FedRecs. It is important to note that in this paper, our focus is on providing a framework to achieve model heterogeneity in FedRecs, while the optimal client division estimation can be explored in future research.

I. Impact of Model Size (RQ5)

Since this paper aims to propose a personalized model sizes federated recommender system, the settings of model sizes are important to be explored. In the main experiments, we set $\{N_s, N_m, N_l\}$ to $\{8, 16, 32\}$. In this part, we tried a smaller setting (i.e., $\{2, 4, 8\}$) and a larger setting (i.e., $\{32, 64, 128\}$). We only present the experimental results on ML in Table VII due to space limitation. A similar conclusion can be obtained from other datasets. At first, by comparing the homogeneous FedRecs (“All Small” and “All Large”), we observed that

TABLE VI: The performance of HeteFedRec under different client group division. “x:y:z” represents $\frac{x}{x+y+z}$, $\frac{y}{x+y+z}$, and $\frac{z}{x+y+z}$ proportions of users are classified into $\mathcal{U}_s, \mathcal{U}_m$, and \mathcal{U}_l according to interaction amounts, respectively. From 5:3:2 to 2:3:5, more and more clients are assigned larger models. The best values on each dataset are bold.

	Datasets	Metrics	All Small	5:3:2	1:1:1	2:3:5	All Large
Fed-NCF	ML	Recall	0.02203	0.02662	0.02619	0.02576	0.02558
		NDCG	0.04328	0.04781	0.04362	0.04197	0.04028
	Anime	Recall	0.04301	0.05855	0.05792	0.05602	0.02727
		NDCG	0.04962	0.05655	0.05540	0.05431	0.04442
	Douban	Recall	0.00759	0.01101	0.00819	0.00756	0.00726
		NDCG	0.01087	0.01290	0.00996	0.00893	0.00878
Fed-LightGCN	ML	Recall	0.02251	0.02434	0.02335	0.02296	0.02301
		NDCG	0.04232	0.04313	0.04233	0.04204	0.04197
	Anime	Recall	0.02924	0.03306	0.03233	0.03064	0.02825
		NDCG	0.04824	0.05177	0.04838	0.04804	0.04788
	Douban	Recall	0.00350	0.00393	0.00360	0.00333	0.00234
		NDCG	0.00530	0.00639	0.00528	0.00483	0.00378

TABLE VII: The performance (NDCG@20) of HeteFedRec and baselines under different model size settings on ML. A similar trend can be observed in Anime and Douban. $\{a, b, c\}$ represents the size of N_s, N_m and N_l . From the left to the right part of the table, model sizes are increased. The best values are bold.

		{2,4,8}	{8,16,32}	{32,64,128}
Fed-NCF	All Small	0.03791	0.04328	0.04028
	All Large	0.04328	0.04028	0.03903
	HeteFedRec	0.03829	0.04781	0.04074
Fed-LightGCN	All Small	0.03813	0.04232	0.04197
	All Large	0.04232	0.04197	0.03901
	HeteFedRec	0.04017	0.04313	0.04093

as the model size increased from 2 to 128, the FedRec’s performance initially improved and then decreased for both Fed-NCF and Fed-LightGCN. For example, in Fed-NCF, when N_s is 2, the NDCG score is 0.03791, indicating that such a small model cannot capture the complex patterns of user-item interactions. Then, when the model size increased to 8, the performance improved to approximately 0.043. However, when the model size continued to increase to 32 and 128, FedRecs’ performance dropped to about 0.040 NDCG scores.

Comparing HeteFedRec with the homogeneous FedRecs, when the setting is $\{2, 4, 8\}$, its performance is better than “All Small” but is slightly behind “All Large”. This is because, in such a range, the models are too small, and simply increasing the model size can yield positive feedback. Under the setting of $\{8, 16, 32\}$, HeteFedRec outperforms both “All Small” and “All Large” due to its personalized model size assignment. For $\{32, 64, 128\}$, where the model sizes are too large, “All Small” achieves the best performance, but our HeteFedRec still outperforms “All Large”. In this paper, our focus is on providing a FedRec framework that enables heterogeneous model sizes. The task of finding an optimal model size set can be explored in future work.

J. Impact of Regularization Factor α (RQ6)

In this part, we explore the impacts of the dimensional decorrelation regularization factor α . Fig. 8 shows the performance changes as α increases from 0.5 to 2.0 on ML due to space limitation. In both cases, the trends indicate that




Fig. 8: The performance trend w.r.t. the change of α for Fed-NCF and Fed-LightGCN on ML. Similar trends can be seen on the other two datasets.

HeteFedRec’s performance initially increases to a peak point and then decreases with a further increase in α .

VI. CONCLUSION

This paper introduces HeteFedRec, the first federated recommendation framework that allows personalized model sizing for participants. HeteFedRec leverages unified dual-task learning to facilitate additive aggregation of recommendation models with varying item embedding sizes. A dimensional decorrelation regularization is employed to prevent dimensional collapse. Additionally, HeteFedRec incorporates a relation-based knowledge distillation approach to enhance knowledge sharing from diverse recommendation models. Extensive experiments are conducted on three real-world datasets using two widely used recommender systems. The results demonstrate the effectiveness and generalizability of HeteFedRec. However, HeteFedRec still requires the recommendation model to have the same model architecture. Besides, how to find the optimal solution of client group division and model sizes for each group is also non-trivial as HeteFedRec’s performance is very sensitive to these settings. In future work, we would like to explore address these two problems.

ACKNOWLEDGMENT

This work is partially supported by the Australian Research Council under the streams of Future Fellowship (No. FT210100624) and Discovery Project (DP240101108).

REFERENCES

- [1] T. Chen, H. Yin, G. Ye, Z. Huang, Y. Wang, and M. Wang, "Try this instead: Personalized and interpretable substitute recommendation," in *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, 2020, pp. 891–900.
- [2] K. Wei, J. Huang, and S. Fu, "A survey of e-commerce recommender systems," in *2007 international conference on service systems and service management*. IEEE, 2007, pp. 1–5.
- [3] H. Yin and B. Cui, *Spatio-temporal recommendation in social media*. Springer, 2016.
- [4] F. Wu, Y. Qiao, J.-H. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu *et al.*, "Mind: A large-scale dataset for news recommendation," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 3597–3606.
- [5] B. Hao, J. Zhang, H. Yin, C. Li, and H. Chen, "Pre-training graph neural networks for cold-start users and items representation," in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 2021, pp. 265–273.
- [6] T. Chen, H. Yin, H. Chen, L. Wu, H. Wang, X. Zhou, and X. Li, "Tada: trend alignment with dual-attention multi-task recurrent neural networks for sales prediction," in *2018 IEEE international conference on data mining (ICDM)*. IEEE, 2018, pp. 49–58.
- [7] H. Yin, B. Cui, Z. Huang, W. Wang, X. Wu, and X. Zhou, "Joint modeling of users' interests and mobility patterns for point-of-interest recommendation," in *Proceedings of the 23rd ACM international conference on Multimedia*, 2015, pp. 819–822.
- [8] S. Zhang, L. Yao, A. Sun, and Y. Tay, "Deep learning based recommender system: A survey and new perspectives," *ACM computing surveys (CSUR)*, vol. 52, no. 1, pp. 1–38, 2019.
- [9] F. Liang, W. Pan, and Z. Ming, "Fedrec++: Lossless federated recommendation with explicit feedback," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 5, 2021, pp. 4224–4231.
- [10] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [11] L. Li, Y. Fan, M. Tse, and K.-Y. Lin, "A review of applications in federated learning," *Computers & Industrial Engineering*, vol. 149, p. 106854, 2020.
- [12] M. Ammad-Ud-Din, E. Ivannikova, S. A. Khan, W. Oyomno, Q. Fu, K. E. Tan, and A. Flanagan, "Federated collaborative filtering for privacy-preserving personalized recommendation system," *arXiv preprint arXiv:1901.09888*, 2019.
- [13] K. Muhammad, Q. Wang, D. O'Reilly-Morgan, E. Tragos, B. Smyth, N. Hurley, J. Geraci, and A. Lawlor, "Fedfast: Going beyond average for faster training of federated recommender systems," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 1234–1242.
- [14] C. Wu, F. Wu, Y. Cao, Y. Huang, and X. Xie, "Fedgnn: Federated graph neural network for privacy-preserving recommendation," *arXiv preprint arXiv:2102.04925*, 2021.
- [15] Z. Sun, Y. Xu, Y. Liu, W. He, Y. Jiang, F. Wu, and L. Cui, "A survey on federated recommendation systems," *arXiv preprint arXiv:2301.00767*, 2022.
- [16] M. Imran, H. Yin, T. Chen, Q. V. H. Nguyen, A. Zhou, and K. Zheng, "Refsr: Resource-efficient federated recommender system for dynamic and diversified user preferences," *ACM Transactions on Information Systems*, vol. 41, no. 3, pp. 1–30, 2023.
- [17] F. M. Harper and J. A. Konstan, "The movielens datasets: History and context," *Acm transactions on interactive intelligent systems (tiis)*, vol. 5, no. 4, pp. 1–19, 2015.
- [18] A. Z. Tan, H. Yu, L. Cui, and Q. Yang, "Towards personalized federated learning," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [19] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016.
- [20] X. Ma, J. Zhu, Z. Lin, S. Chen, and Y. Qin, "A state-of-the-art survey on solving non-iid data in federated learning," *Future Generation Computer Systems*, vol. 135, pp. 244–258, 2022.
- [21] H. Chang, V. Shejwalkar, R. Shokri, and A. Houmansadr, "Cronus: Robust and heterogeneous collaborative learning with black-box knowledge transfer," *arXiv preprint arXiv:1912.11279*, 2019.
- [22] D. Li and J. Wang, "Fedmd: Heterogenous federated learning via model distillation," *arXiv preprint arXiv:1910.03581*, 2019.
- [23] L. Sun and L. Lyu, "Federated model distillation with noise-free differential privacy," *arXiv preprint arXiv:2009.05537*, 2020.
- [24] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [25] E. Diao, J. Ding, and V. Tarokh, "Heteroff: Computation and communication efficient federated learning for heterogeneous clients," *arXiv preprint arXiv:2010.01264*, 2020.
- [26] Z. Zhu, J. Hong, S. Drew, and J. Zhou, "Resilient and communication efficient learning for heterogeneous federated systems," in *Proceedings of Thirty-ninth International Conference on Machine Learning (ICML 2022)*, 2022.
- [27] K. Wang, Q. He, F. Chen, C. Chen, F. Huang, H. Jin, and Y. Yang, "Flexified: Personalized federated learning for edge clients with heterogeneous model architectures," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 2979–2990.
- [28] R. Liu, F. Wu, C. Wu, Y. Wang, L. Lyu, H. Chen, and X. Xie, "No one left behind: Inclusive federated learning over heterogeneous devices," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 3398–3406.
- [29] M. Kim, S. Yu, S. Kim, and S.-M. Moon, "Depthfl: Depthwise federated learning for heterogeneous clients," in *The Eleventh International Conference on Learning Representations*, 2023.
- [30] Y. Li, T. Chen, P.-F. Zhang, and H. Yin, "Lightweight self-attentive sequential recommendation," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 967–977.
- [31] N. Q. V. Hung, H. H. Viet, N. T. Tam, M. Weidlich, H. Yin, and X. Zhou, "Computing crowd consensus with partial agreement," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 1, pp. 1–14, 2017.
- [32] Q. V. H. Nguyen, C. T. Duong, T. T. Nguyen, M. Weidlich, K. Aberer, H. Yin, and X. Zhou, "Argument discovery via crowdsourcing," *The VLDB Journal*, vol. 26, pp. 511–535, 2017.
- [33] J. Yu, H. Yin, X. Xia, T. Chen, J. Li, and Z. Huang, "Self-supervised learning for recommender systems: A survey," *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [34] R. Zheng, L. Qu, B. Cui, Y. Shi, and H. Yin, "Automl for deep recommender systems: A survey," *ACM Transactions on Information Systems*, vol. 41, no. 4, pp. 1–38, 2023.
- [35] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.
- [36] T. Qi, F. Wu, C. Wu, Y. Huang, and X. Xie, "Privacy-preserving news recommendation model learning," *arXiv preprint arXiv:2003.09592*, 2020.
- [37] Z. Liu, L. Yang, Z. Fan, H. Peng, and P. S. Yu, "Federated social recommendation with graph neural network," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 13, no. 4, pp. 1–24, 2022.
- [38] J. Long, T. Chen, Q. V. H. Nguyen, G. Xu, K. Zheng, and H. Yin, "Model-agnostic decentralized collaborative learning for on-device poi recommendation," in *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2023, pp. 423–432.
- [39] Q. Wang, H. Yin, T. Chen, J. Yu, A. Zhou, and X. Zhang, "Fast-adapting and privacy-preserving federated recommender system," *The VLDB Journal*, pp. 1–20, 2021.
- [40] C. Wu, F. Wu, T. Qi, Y. Huang, and X. Xie, "Fedcl: Federated contrastive learning for privacy-preserving recommendation," *arXiv preprint arXiv:2204.09850*, 2022.
- [41] F. Chen, M. Luo, Z. Dong, Z. Li, and X. He, "Federated meta-learning with fast convergence and efficient communication," *arXiv preprint arXiv:1802.07876*, 2018.
- [42] H. Zhang, F. Luo, J. Wu, X. He, and Y. Li, "Lightfr: Lightweight federated recommendation with privacy-preserving matrix factorization," *ACM Transactions on Information Systems*, vol. 41, no. 4, pp. 1–28, 2023.
- [43] S. Luo, Y. Xiao, and L. Song, "Personalized federated recommendation via joint representation learning, user clustering, and model adaptation," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022, pp. 4289–4293.

- [44] S. Zhang, H. Yin, T. Chen, Z. Huang, Q. V. H. Nguyen, and L. Cui, "Pipattack: Poisoning federated recommender systems for manipulating item promotion," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 2022, pp. 1415–1423.
- [45] D. Rong, S. Ye, R. Zhao, H. N. Yuen, J. Chen, and Q. He, "Fedrecattack: Model poisoning attack to federated recommendation," in *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2022, pp. 2643–2655.
- [46] W. Yuan, Q. V. H. Nguyen, T. He, L. Chen, and H. Yin, "Manipulating federated recommender systems: Poisoning with synthetic users and its countermeasures," *arXiv preprint arXiv:2304.03054*, 2023.
- [47] W. Yuan, S. Yuan, C. Yang, Q. V. H. Nguyen, and H. Yin, "Manipulating visually-aware federated recommender systems and its countermeasures," *ACM Transactions on Information Systems*, 2023.
- [48] S. Zhang and H. Yin, "Comprehensive privacy analysis on federated recommender system against attribute inference attacks," *arXiv preprint arXiv:2205.11857*, 2022.
- [49] W. Yuan, C. Yang, Q. V. H. Nguyen, L. Cui, T. He, and H. Yin, "Interaction-level membership inference attack against federated recommender systems," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 1053–1062.
- [50] W. Yuan, H. Yin, F. Wu, S. Zhang, T. He, and H. Wang, "Federated unlearning for on-device recommendation," in *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 2023, pp. 393–401.
- [51] S. Zhang, H. Yin, T. Chen, Z. Huang, L. Cui, and X. Zhang, "Graph embedding for recommendation against attribute inference attacks," in *Proceedings of the Web Conference 2021*, 2021, pp. 3002–3014.
- [52] D. Gao, X. Yao, and Q. Yang, "A survey on heterogeneous federated learning," *arXiv preprint arXiv:2210.04505*, 2022.
- [53] X. Sun, H. Yin, B. Liu, H. Chen, J. Cao, Y. Shao, and N. Q. Viet Hung, "Heterogeneous hypergraph embedding for graph classification," in *Proceedings of the 14th ACM international conference on web search and data mining*, 2021, pp. 725–733.
- [54] G. Ye, T. Chen, Y. Li, L. Cui, Q. V. H. Nguyen, and H. Yin, "Heterogeneous collaborative learning for personalized healthcare analytics via messenger distillation," *IEEE Journal of Biomedical and Health Informatics*, 2023.
- [55] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.
- [56] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.
- [57] A. Li, J. Sun, P. Li, Y. Pu, H. Li, and Y. Chen, "Hermes: an efficient federated learning framework for heterogeneous mobile clients," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 2021, pp. 420–437.
- [58] Q. Li, B. He, and D. Song, "Model-contrastive federated learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 713–10 722.
- [59] Z. Jiang, Y. Xu, H. Xu, Z. Wang, C. Qiao, and Y. Zhao, "Fedmp: Federated learning through adaptive model pruning in heterogeneous edge computing," in *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2022, pp. 767–779.
- [60] H. Wang, S. Marella, and J. Anderson, "Fedadmm: A federated primal-dual algorithm allowing partial participation," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 287–294.
- [61] S. Horvath, S. Laskaridis, M. Almeida, I. Leontiadis, S. Venieris, and N. Lane, "Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 876–12 889, 2021.
- [62] L. Yuan, F. E. Tay, G. Li, T. Wang, and J. Feng, "Revisiting knowledge distillation via label smoothing regularization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3903–3911.
- [63] J. Tang, R. Shivanna, Z. Zhao, D. Lin, A. Singh, E. H. Chi, and S. Jain, "Understanding and improving knowledge distillation," *arXiv preprint arXiv:2002.03532*, 2020.
- [64] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," *arXiv preprint arXiv:1612.03928*, 2016.
- [65] P. Chen, S. Liu, H. Zhao, and J. Jia, "Distilling knowledge via knowledge review," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5008–5017.
- [66] Y. Shang, B. Duan, Z. Zong, L. Nie, and Y. Yan, "Lipschitz continuity guided knowledge distillation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 675–10 684.
- [67] C. Yang, Z. An, L. Cai, and Y. Xu, "Mutual contrastive learning for visual representation learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 3045–3053.
- [68] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 173–182.
- [69] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "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*, 2020, pp. 639–648.
- [70] T. Hua, W. Wang, Z. Xue, S. Ren, Y. Wang, and H. Zhao, "On feature decorrelation in self-supervised learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9598–9608.
- [71] Y. Shi, J. Liang, W. Zhang, V. Y. Tan, and S. Bai, "Towards understanding and mitigating dimensional collapse in heterogeneous federated learning," *arXiv preprint arXiv:2210.00226*, 2022.
- [72] F. Zhu, C. Chen, Y. Wang, G. Liu, and X. Zheng, "Dtcd: A framework for dual-target cross-domain recommendation," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 1533–1542.
- [73] X. Ren, W. Yuan, T. Chen, C. Yang, Q. V. H. Nguyen, and H. Yin, "Joint semantic and structural representation learning for enhancing user preference modelling," *arXiv preprint arXiv:2304.12083*, 2023.
- [74] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, "An efficient framework for clustered federated learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 19 586–19 597, 2020.
- [75] Q. Meng, F. Zhou, H. Ren, T. Feng, G. Liu, and Y. Lin, "Improving federated learning face recognition via privacy-agnostic clusters," *arXiv preprint arXiv:2201.12467*, 2022.
- [76] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.