

Uni-FedRec: A Unified Privacy-Preserving News Recommendation Framework for Model Training and Online Serving

Tao Qi¹, Fangzhao Wu², Chuhan Wu¹, Yongfeng Huang¹ and Xing Xie²

¹Department of Electronic Engineering & BNRist, Tsinghua University, Beijing 100084, China

²Microsoft Research Asia, Beijing 100080, China

{taoqi.qt, wufangzhao, wuchuhan15}@gmail.com

yfhuang@tsinghua.edu.cn

xing.xie@microsoft.com

Abstract

News recommendation is important for personalized online news services. Most existing news recommendation methods rely on centrally stored user behavior data to both train models offline and provide online recommendation services. However, user data is usually highly privacy-sensitive, and centrally storing them may raise privacy concerns and risks. In this paper, we propose a unified news recommendation framework, which can utilize user data locally stored in user clients to train models and serve users in a privacy-preserving way. Following a widely used paradigm in real-world recommender systems, our framework contains two stages. The first one is for candidate news generation (i.e., recall) and the second one is for candidate news ranking (i.e., ranking). At the recall stage, each client locally learns multiple interest representations from clicked news to comprehensively model user interests. These representations are uploaded to the server to recall candidate news from a large news pool, which are further distributed to the user client at the ranking stage for personalized news display. In addition, we propose an interest decomposer-aggregator method with perturbation noise to better protect private user information encoded in user interest representations. Besides, we collaboratively train both recall and ranking models on the data decentralized in a large number of user clients in a privacy-preserving way. Experiments on two real-world news datasets show that our method can outperform baseline methods and effectively protect user privacy.

1 Introduction

Online news platforms usually rely on personalized news recommendation techniques to help users obtain their interested news information (Qi et al., 2021b; Wu et al., 2019c). Existing news recommendation models usually exploit users' historical behavior data to model user interests for matching candidate news (Wang et al., 2020; Wu et al.,

2019c,b, 2020c; Qi et al., 2021a; Ge et al., 2020; Wu et al., 2021c). For example, Okura et al. (2017) employed a GRU network to build user embeddings from browsed news. Wu et al. (2019a) employed an attention network to build user embeddings by aggregating different clicked news. Both of them match candidate news and user interests via the inner product of their embeddings. In short, most of these methods rely on centralized storage of user behavior data to train models and serve users. However, user behavior data is usually highly privacy-sensitive (Chai et al., 2019), and centrally storing them may arouse users' concerns on privacy leakage and violate some privacy protection regulations such as GDPR¹.

A few methods explore to recommend news in a privacy-preserving way (Qi et al., 2020). For instance, Qi et al. (2020) proposed to store user data in user clients and applied federated learning technique (McMahan et al., 2017) to train news recommendation models on decentralized data. In general, these methods usually focus on developing privacy-preserving model training approaches based on decentralized user behavior data for ranking candidate news. However, how to generate candidate news and serve users in a privacy-preserving way remains an open problem.

In this paper, we propose a unified news recommendation framework based on federated learning (named *Uni-FedRec*), which can utilize user behavior data locally stored in user clients to train models offline and serve users online in a privacy-preserving way. Following a widely applied paradigm in real-world recommender systems (Wu et al., 2021e; Pal et al., 2020), *Uni-FedRec* contains a recall stage for personalized candidate news generation and a ranking stage for candidate news ranking. In the recall stage, the user client first locally learns multiple interest representations from clicked news to model diverse

¹<https://gdpr-info.eu/>

user interests. These representations are further uploaded to the server to recall a small number of candidate news (e.g., 100) from a large news pool. In the ranking stage, recalled candidate news are distributed to the user client and locally ranked for news personalized display. Besides, user interest representations may encode user privacy information. (Wu et al., 2020a). To protect private user information encoded in interest representations, we propose an interest decomposer-aggregator method with perturbation noise to synthesize interest representations with a group of basic interest embeddings. In addition, *Uni-FedRec* utilizes user data decentralized in a large number of user clients to collaboratively train the recall and ranking model in a privacy-preserving way. Extensive experiments on two real-world datasets verify that our method can significantly outperform baseline methods and effectively protect user privacy.

In summary, our contributions are as follows:

- We propose a unified privacy-persevering news recommendation framework which can train model offline and serve users online with locally stored user data.
- We propose a privacy-preserving recall model which can comprehensively model diverse user interests and protect user privacy.
- Extensive experiments demonstrate that our framework can outperform many baseline methods and effectively protect user privacy.

2 Related Work

2.1 Personalized News Recommendation

Personalized news recommendation is an important research problem and has been widely studied in recent years (Konstan et al., 1997; Wang and Blei, 2011; Liu et al., 2010; Bansal et al., 2015; Wu et al., 2020b; Qi et al., 2021c; Wu et al., 2020d, 2021d; Wang et al., 2020; Ge et al., 2020; An et al., 2019). Existing news recommendation methods aim to match candidate news content with user preferences mined from users’ historical behaviors (Khattar et al., 2018; Wu et al., 2021b,f; Ge et al., 2020; Qi et al., 2021a; Wu et al., 2019d; An et al., 2019). For example, Okura et al. (2017) proposed to learn user interest embeddings from the sequential information of user’s clicked news via a GRU network. An et al. (2019) proposed to model short-term user interest from news clicks via a GRU network and

model long-term user interest via user ID embeddings. They further combine them to form a unified interest embedding. Wu et al. (2019d) employed a multi-head self-attention network to learn user interest embeddings by modeling relatedness of users’ reading behaviors. Besides, all of these three methods performed the matching between candidate news and user interest via the inner product of their embeddings. In brief, most of these methods rely on the centralized storage of user behavior data to train models and serve users. However, users’ behavior data is usually highly privacy-sensitive, and storing them in the server may arouse risks and user concerns on privacy leakage, and may also violate some privacy protection regulations (e.g., GDPR) (Muhammad et al., 2020; Wu et al., 2020a). Different from these methods, we propose a unified privacy-preserving framework for news recommendation, which can utilize decentralized user behavior data to train models and serve users.

2.2 Privacy-Preserving Recommendation

Recently, due to users’ increasing concerns on privacy leakage, some privacy-preserving recommendation methods have been proposed (Qi et al., 2020; Flanagan et al., 2020; Lin et al., 2020; Wang et al., 2021; Muhammad et al., 2020; Yang et al., 2021; Wu et al., 2021a, 2020a). For example, Chai et al. (2019) proposed to compute gradients of user and item embeddings in user clients based on locally stored user rating data and upload gradients to the server for federated model updating. Besides, to better protect user privacy, they employed the homomorphic encryption technique (Gentry, 2009) to encrypt the uploaded gradients. Qi et al. (2020) proposed to apply federated learning technique to train neural news recommendation models on decentralized user data. They used local differential privacy technique (Ren et al., 2018) to protect the uploaded gradients from leaking user privacy. In brief, most of these methods focus on training a recommendation model for ranking candidate news in a privacy-preserving way. However, how to generate candidate news from news pool according to user interest and serve users with decentralized user behavior data are still unsolved problems. Different from these methods, we propose a unified privacy-preserving news recommendation framework, which can utilize locally stored user data to generate candidate news from the server, and further serve users via local candidate news ranking.

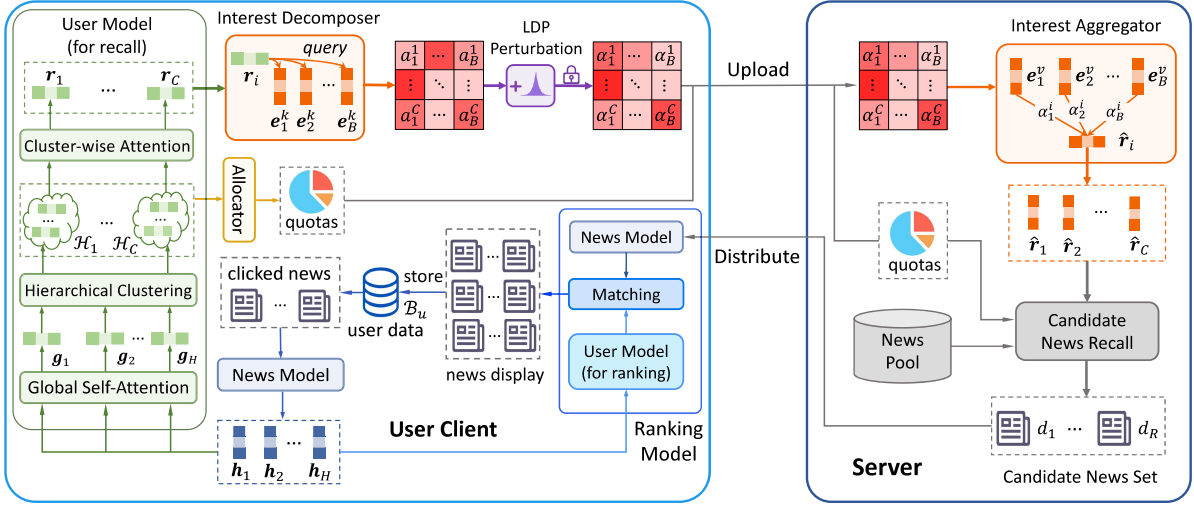


Figure 1: The framework of *Uni-FedRec* for privacy-preserving online serving.

3 Uni-FedRec

In this section, we will introduce our unified privacy-preserving news recommendation framework (named *Uni-FedRec*), which can utilize decentralized user data to train models and serve users.

3.1 Framework Overview

In *Uni-FedRec*, user behavior data (e.g., displayed news and clicked news) is locally stored in user clients, and the news pool is stored and maintained in the server. Following a widely used paradigm in real-world recommender systems (Pal et al., 2020; Liu et al., 2020), *Uni-FedRec* contains a recall stage for candidate news generation and a ranking stage for candidate news ranking. To serve a user, the user client first employs a privacy-preserving recall model to locally learn multiple interest representations from clicked news to model diverse user interests. The interest representations are further uploaded to the server to recall candidate news from a large news pool. In the ranking stage, recalled candidate news are distributed to the user client and locally ranked for personalized news display. To train models on decentralized user data, *Uni-FedRec* coordinates massive user clients to collaboratively calculate gradients from their local user data for federated model updating. Next, we will introduce each module of *Uni-FedRec* in detail.

3.2 Privacy-Preserving Recall Model

As shown in Fig.1, our privacy-preserving recall model contains four major modules, i.e., a *user model*, an *interest decomposer*, an *LDP perturbation module* and an *interest aggregator*. The former

is used to learn multiple interest representations to model diverse user interests. The latter three are used to protect users' private information encoded in interest representations.

Interest Modeling: User behavior contexts are informative for modeling user interests (Wu et al., 2019d). In *user model*, we first use a global self-attention network (Vaswani et al., 2017) to learn context-sensitive representations of clicked news $[\mathbf{g}_1, \dots, \mathbf{g}_H]$ from representations of clicked news $[\mathbf{h}_1, \dots, \mathbf{h}_H]$, where H is the number of clicked news². Besides, users usually have diverse interests in multiple interest fields (Pal et al., 2020; Liu et al., 2020). To capture diverse user interests, we divide user's clicked news into different interest clusters $\{\mathcal{H}_i | i = 1, \dots, C\}$ via the hierarchical clustering algorithm (Johnson, 1967), where \mathcal{H}_i is the i -th cluster, and C is the number of clusters. The algorithm hierarchically merges the clusters until the average distances between any two clusters are larger than a given threshold d_c . Then, we apply a cluster-wise attention network to learn unified interest representation \mathbf{r}_i for each cluster \mathcal{H}_i :

$$\mathbf{r}_i = \sum_{j=1}^{|\mathcal{H}_i|} \gamma_j^i \mathbf{g}_j^i, \quad \gamma_j^i = \frac{\exp(\text{Att}(\mathbf{g}_j^i))}{\sum_{k=1}^{|\mathcal{H}_i|} \exp(\text{Att}(\mathbf{g}_k^i))}, \quad (1)$$

where γ_j^i and \mathbf{g}_j^i is the attention weight and representation for the j -th clicked news in the i -th cluster, and $\text{Att}(\cdot)$ is a dense network for calculating attention scores. In this way, we can obtain multiple interest representations $\{\mathbf{r}_i | i = 1, \dots, C\}$ to model user interests in different interest fields.

²We introduced how to learn news representations from news texts in Section 3.3.

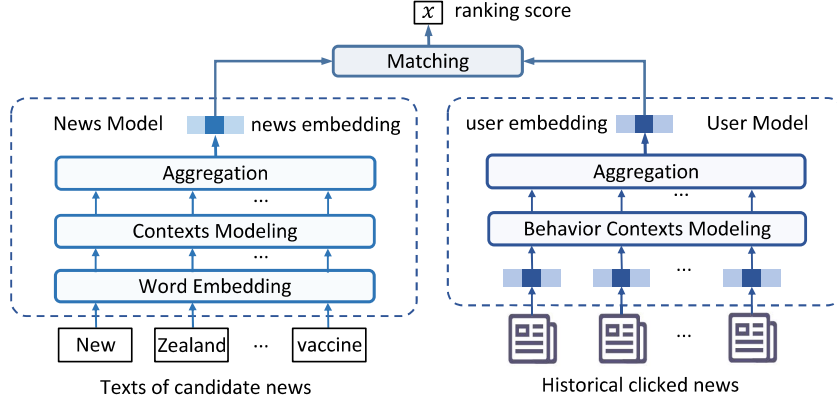


Figure 2: Framework of the news ranking model.

Privacy Protection: User interest representations may contain some private user information (Wu et al., 2020a). To protect private information encoded in interest representations, we propose an interest aggregator-decomposer method with permutation noise. Its core is to synthesize interest representations by combining a group of trainable basic interest embeddings $\{(\mathbf{e}_i^k, \mathbf{e}_i^v) | i = 1, \dots, B\}$ (*BIE*) shared among different users, where \mathbf{e}_i^k and \mathbf{e}_i^v is the key and value of the i -th basic interest embedding, respectively, and B is the number of *BIE*. In *interest decomposer*, we first decompose each interest representation \mathbf{r}_i on the keys of *BIE*: $a_j^i = \text{Query}(\mathbf{r}_i, \mathbf{e}_j^k)$, where $\text{Query}(\mathbf{x}, \mathbf{y})$ is a query function implemented by dot product of \mathbf{x} and \mathbf{y} , $a_j^i \in \mathbb{R}$ is the decomposition score of decomposing \mathbf{r}_i on embedding \mathbf{e}_j^k . We further perturb decomposition scores via local differential privacy (LDP) technique:

$$\hat{a}_j^i = f_\delta(a_j^i) + n_I, \quad n_I \sim \text{La}(0, \lambda_I), \quad (2)$$

where \hat{a}_j^i is the protected decomposition score, $f_\delta(z)$ is a function for clipping z with the scale of δ , n_I is a zero-mean Laplace noise, and λ_I is its intensity. Next, in the *interest aggregator*, we further synthesize protected interest representations by combining value embeddings of *BIE*:

$$\hat{\mathbf{r}}_i = \sum_{j=1}^B \alpha_j^i \mathbf{e}_j^v, \quad \alpha_j^i = \frac{\exp(\hat{a}_j^i)}{\sum_{k=1}^B \exp(\hat{a}_k^i)}, \quad (3)$$

where $\hat{\mathbf{r}}_i$ is the protected representation for \mathbf{r}_i .

News Recall: We further use each protected interest representation $\hat{\mathbf{r}}_i$ to recall top R_i candidate news that has the largest relevance with $\hat{\mathbf{r}}_i$ from news in the pool. We use inner product similarity to measure representation relevance, and this recall progress can be speeded up by some search

algorithms such as ANN search (Arya et al., 1998). Besides, we have an *allocator* to allocate quotas, i.e., the number of candidate news recalled by each channel. We utilize ratios of clicked news belonging to each interest channel to generate their quotas:

$$R_i = q_i \times R, \quad q_i = \frac{|\mathcal{H}_i|}{\sum_{j=1}^C |\mathcal{H}_j|} \quad (4)$$

where R is the total number of recalled candidate news. Finally, we integrated candidate news of different channels and obtain the candidate news set $\mathcal{R} = \{d_i | i = 1, \dots, R\}$, where d_i is the i -th recalled candidate news.

Loss Function: InfoNCE loss (Oord et al., 2018; Wu et al., 2019d; An et al., 2019) is usually used to formulate loss function in recommendation task. It requires a unified score to rank positive and negative samples, while our method will generate C different recall scores for each news. To tackle this issue, we combine recall scores of news d generated by each channel \hat{z}_i to form a unified score:

$$z = \sum_{i=1}^C q_i \hat{z}_i, \quad \hat{z}_i = \mathbf{d} \cdot \hat{\mathbf{r}}_i, \quad (5)$$

where \mathbf{d} is representation of d . Next, for each positive sample, we randomly select K_r negative samples from all news that are displayed to but not clicked by this user. Then, we obtain the loss function \mathcal{L}_u^r based on behavior data \mathcal{B}_u of user u :

$$\mathcal{L}_u^r = \sum_{i=1}^{|\mathcal{B}_u|} \frac{\exp(z_i)}{\exp(z_i) + \sum_{j=1}^{K_r} \exp(z_i^j)}, \quad (6)$$

where z_i and z_i^j is the unified score of the i -th positive sample and its j -th negative sample.

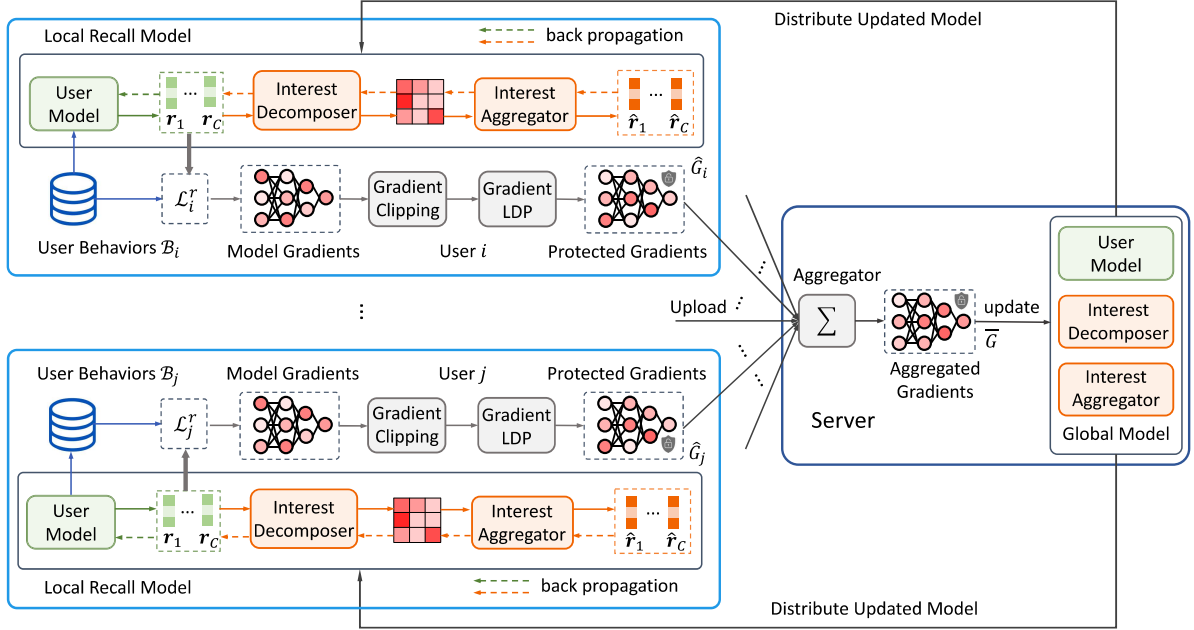


Figure 3: The framework of *Uni-FedRec* for privacy-preserving model training.

3.3 Ranking Model

Uni-FedRec contains a ranking model to locally rank candidate news in the user client. Since local news ranking will not leak user privacy, we directly employ existing news ranking methods such as *NRMS* (Wu et al., 2019d) in *Uni-FedRec*. As shown in Fig. 2, these methods share a similar framework, where a news model learns news embedding from news texts, a user model learns user embedding from clicked news, and a matching module (e.g., dot product) matches candidate news and user interests for personalized ranking.³ The news model is usually based on the stack of a word embedding layer, a context modeling layer (e.g., Transformer) and an aggregation layer (e.g., attention network), and the user model is usually based on the stack of a behavior context modeling layer (e.g., GRU) and an aggregation layer (e.g., attention network). We also formulate the loss function \mathcal{L}_u^g of the ranking model via the infoNCE loss:

$$\mathcal{L}_u^g = \frac{1}{|\mathcal{B}_u|} \sum_{i=1}^{|\mathcal{B}_u|} \frac{\exp(x_i)}{\exp(x_i) + \sum_{j=1}^{K_g} \exp(x_i^j)}, \quad (7)$$

where x_i and x_i^j is the ranking score of the i -th positive sample and its j -th negative sample randomly selected from the same news impression respectively and K_g is the number of negative samples.

³We utilize news model in the ranking model to generate news representation for the recall model.

3.4 Privacy-Preserving Online Serving

In Fig. 1, we show the privacy-preserving framework of *Uni-FedRec* for online serving with decentralized user data. In *Uni-FedRec*, users' behavior data is locally stored in user clients and is never uploaded to the server, which can effectively alleviate users' privacy concerns. The server stores and maintains news articles in a news pool. Besides, both the user client and the server contain the whole recall model. When a user visits the news platform, the client first employs the recall model (i.e., *user model*, *interest decomposer* and *perturbation module*) to build protected decomposition weights α_j^i and generate quotas of interest channels, which are further uploaded to the server.⁴ After receiving them, the server further employs the *interest aggregator* to build protected interest representations $\hat{\mathbf{r}}_i$ and perform news recall to obtain candidate news set \mathcal{R} . The recalled news with their titles and links are further distributed to the user client. After the client receives the candidate news, the ranking model locally ranks these candidate news and displays the top D news with the highest-ranking scores. The user can locally read news titles and click links of interested news for reading. Finally, after the user finishes this session, all displayed news and user behaviors will be stored in the user client. Besides, to reduce online

⁴We upload protected decomposition weights instead of interest representations to solve communication costs.

latency, protected decomposition weights and interest channel quotas can be calculated and stored in advance for quick uploading.

3.5 Privacy-Preserving Model Training

Training accurate recommendation models usually relies on large-scale training data (Muhammad et al., 2020). In *Uni-FedRec*, user behavior data is locally stored in user clients and it is hard to centrally learn parameters of recommendation models in a conventional way. Motivated by McMahan et al. (2017), we utilize federated learning techniques to train recall and ranking models on decentralized user data. Next, we will introduce the training process of the recall model in detail.⁵

As shown in Fig. 3, the server is in charge of maintaining parameters Θ_r of the recall model and coordinating clients to collaboratively update Θ_r based on user data locally stored on them. At the beginning of a federated updating round, each client has a copy of the current recall model. Then the server will randomly select a part of users \mathcal{U} to perform local gradient calculation. Client of the selected user u will calculate gradients $\frac{\partial \mathcal{L}_u^r}{\partial \Theta_r}$ for parameter updating based on her behavior data \mathcal{B}_u and current model parameters Θ_r . Besides, gradients calculated from user behavior data may leak private user information (Zhu et al., 2019). To better protect user privacy, following Qi et al. (2020), we apply a local differential privacy (LDP) module to the clipped gradients:

$$\hat{\mathbf{G}}_u = f_\theta\left(\frac{\partial \mathcal{L}_u^r}{\partial \Theta_r}\right) + n_g, \quad n_g \sim La(0, \lambda_g), \quad (8)$$

where $\hat{\mathbf{G}}_u$ is the protected gradients, f_θ is a clipping function with the scale of θ , n_g is the Laplace noise and λ_g is its intensity. Next, the user client uploads the protected gradients $\hat{\mathbf{G}}_u$ to the server.

After the server receives uploaded gradients from clients of users in \mathcal{U} , the server further aggregates these gradients for model updating:

$$\bar{\mathbf{G}} = \sum_{u \in \mathcal{U}} \beta_u \hat{\mathbf{G}}_u, \quad \beta_u = \frac{|\mathcal{B}_u|}{\sum_{v \in \mathcal{U}} |\mathcal{B}_v|}, \quad (9)$$

where $\bar{\mathbf{G}}$ is the average gradients. Then, parameters of the recall model is updated as: $\Theta_r = \Theta_r - \omega \bar{\mathbf{G}}$, where ω is the learning rate. Updated parameters are further distributed to all clients to update local model. We will repeat this federated updating process until the recall model training converges.

⁵Training process of the ranking model is similar to the recall model and its description is omitted.

3.6 Analysis on Privacy Protection

In this section, we will analyze the privacy protection ability of *Uni-FedRec*. First, in *Uni-FedRec*, users' behavior data is locally stored in their clients and is never uploaded to the server, which can effectively alleviate users' privacy concerns and risks of large-scale privacy leakage (McMahan et al., 2017). To train models and serve users, *Uni-FedRec* only needs to upload model gradients and user interest representations to the server. These intermediate variables usually contain much less private information than raw data according to data processing inequality (McMahan et al., 2017; Qi et al., 2020). Besides, these variables are aggregated from multiple behaviors of a user, making it more difficult to infer a specific user behavior from them. Second, we propose an interest decomposer-aggregator method to protect interest representation \mathbf{r}_i . Since protected interest representation $\hat{\mathbf{r}}_i$ is aggregated from basic interest embeddings shared among users instead of user's clicked news, it is more difficult to infer a specific user's clicked news from $\hat{\mathbf{r}}_i$ than \mathbf{r}_i . Besides, in this method, $\mathbf{r}_i \in \mathbb{R}^d$ which belongs to a d -dimensional space \mathbb{R}^d is projected into a B -dimensional space \mathbb{R}^B . Since B is much smaller than d in our settings, $\hat{\mathbf{r}}_i$ can lose much information on user privacy. Third, we apply the LDP technique to protect both interest representations and gradients. Based on the LDP theory (Choi et al., 2018), in *Uni-FedRec*, the privacy budget upper bounds of protected gradients and protected interest representations can achieve $\frac{2\theta}{\lambda_g}$ and $\frac{2\delta}{\lambda_I}$, respectively. Since a smaller privacy budget means better privacy protection ability, *Uni-FedRec* can achieve a trade-off between model accuracy and privacy protection by adjusting noise intensity.

4 Experiment

4.1 Experimental Datasets and Settings

	# News	# Users	# Clicks	# Impressions
<i>MIND</i>	161,013	1,000,000	24,155,470	15,777,377
<i>NewsFeeds</i>	120,219	20,000	112,927	48,923

Table 1: Dataset statistics.

We conduct experiments on two real-world datasets. The first one is *MIND* (Wu et al., 2020e), a public news recommendation dataset constructed by logs of 1 million users in the Microsoft News during six weeks (Oct. 12 to Nov. 22, 2019). The second one is *NewsFeeds*, which is constructed by

	MIND				NewsFeeds			
	R@100	R@200	R@300	R@400	R@100	R@200	R@300	R@400
YoutubeNet	1.50±0.03	2.43±0.08	3.34±0.08	3.96±0.13	0.60±0.02	0.92±0.01	1.17±0.01	1.45±0.02
HUITA	1.69±0.06	2.67±0.04	3.37±0.06	3.97±0.06	0.60±0.01	0.91±0.01	1.18±0.03	1.45±0.01
EBNR	2.31±0.17	3.72±0.13	4.69±0.17	5.61±0.17	0.64±0.03	0.96±0.05	1.28±0.06	1.55±0.06
SASRec	2.22±0.05	3.51±0.07	4.54±0.07	5.38±0.07	0.62±0.06	0.96±0.01	1.20±0.06	1.49±0.05
PinnerSage	1.22±0.14	1.85±0.28	2.69±0.23	3.53±0.20	0.59±0.01	0.93±0.01	1.15±0.01	1.45±0.02
Octopus	1.26±0.03	1.93±0.07	2.74±0.06	3.55±0.06	0.60±0.02	0.92±0.02	1.17±0.02	1.44±0.03
Uni-FedRec	2.95±0.11	4.13±0.12	5.13±0.12	5.99±0.11	0.80±0.08	1.14±0.10	1.60±0.12	2.03±0.12

Table 2: News recall performance of different methods. Higher recall rates mean better performance. T-test on these results verifies the improvement of *Uni-FedRec* over baseline methods is significant at level $p \leq 0.001$.

	MIND				NewsFeeds			
	R@100	R@200	R@300	R@400	R@100	R@200	R@300	R@400
YoutubeNet	12.29	15.91	18.48	20.64	29.43	31.22	32.46	33.47
HUITA	13.44	16.11	17.98	19.49	29.51	31.24	32.44	33.39
EBNR	5.49	8.27	10.30	12.05	11.35	13.08	14.14	14.86
SASRec	6.00	8.71	10.81	12.52	7.78	9.18	10.16	11.03
PinnerSage	16.91	21.35	24.48	27.18	29.43	31.10	32.32	33.38
Octopus	17.04	21.62	24.72	27.31	29.45	31.15	32.36	33.38
Uni-FedRec	0.55	1.14	1.69	2.22	0.23	0.54	0.83	1.08

Table 3: Privacy protection performance of different methods, which is measured by rates of user’s historical clicked news recalled from the news pool. Lower recall rates means better privacy protection performance.

logs of 20,000 users from a commercial news feeds production of Microsoft during two weeks (Mar. 18 to Apr. 1, 2020). User logs in the first week are used to construct historical user behaviors, logs in the last two days are used for evaluation, and other logs are used for model training. More information on *MIND* and *NewsFeeds* is in Table 1.

Next, we will introduce settings of our experiments. In our privacy-preserving news recall model, dimensions of both news and user representations are 256. The self-attention network contains 16 attention heads with 16-dimensional output vectors. The clustering distance d_c is set to 1. The cluster-wise attention network is a two-layer dense network with 128-dimensional hidden vector. The number (B) of basic interest embeddings is set to 30 and dimensions of these basic interest embeddings are 256. The clipping scale δ is set to 0.2 and intensity λ_I of the interest representation perturbation noise n_I is set to 1.2. Besides, we combine four different news ranking models, i.e., *FedRec* (Qi et al., 2020), *NRMS* (Wu et al., 2019d), *LSTUR* (An et al., 2019) and *NAML* (Wu et al., 2019a), with our proposed privacy-preserving news recall model in *Uni-FedRec*. Embeddings generated by the news model and user model in these ranking methods are 256-dimensional. We

randomly sample $r = 2\%$ clients in each round for model updating. The gradient clipping scale θ is 0.1 and intensity λ_g of gradient perturbation noise n_g is 0.01. Negative sampling ratios for training both recall and ranking models, i.e., K_r and K_g , are 4. The learning rate ω is 0.05. Codes of *Uni-FedRec* are released for reproducing our method.⁶

4.2 News Recall Performance

We compare the performance of different recall models on (1) news recall and (2) privacy protection. News recall performance is measured by rates of users’ future clicked news in the top K recalled candidate news ($R@K$). Privacy protection performance is measured by rates of users’ historical clicked news in the top K recalled candidate news. Since it is easier to infer user private information from representations that can recall more users’ historical clicked, methods that achieve lower recall rates on historical clicks are considered to have better privacy protection performance. Here are baseline methods we compared: (1) *YoutubeNet* (Covington et al., 2016): averaging clicked news representations to recall candidate news. (2) *HUITA* (Wu et al., 2019e): attentively aggregating clicked news for news recall. (3) *EBNR* (Okura et al., 2017):

⁶<https://github.com/taoqi98/UniFedRec>

	FedRec	LSTUR	NRMS	NAML
YoutubeNet	70.65	68.53	68.79	65.93
HUITA	70.48	68.76	70.48	68.76
EBNR	75.56	73.82	75.01	70.89
SASRec	75.07	72.51	73.35	70.51
PinnerSage	69.26	68.96	67.28	66.09
Octopus	69.76	69.12	67.11	65.75
Uni-FedRec	79.26	77.31	78.91	75.40

Table 4: Recommendation performance (AUC) of different methods on *MIND*, where rows and columns are different recall and ranking methods, respectively.

learning user representation via a GRU network. (4) *SASRec* (Kang and McAuley, 2018): using a self-attention network to learn user representation. (5) *PinnerSage* (Pal et al., 2020): learning multiple interest representations via news clustering. (6) *Octopus* (Liu et al., 2020): modding multiple user interest via elastic archive network.

We repeat experiment on each method 5 times and show results in Table 2 and Table 3. As shown in Table 2, *Uni-FedRec* significantly outperforms baseline methods in news recall. This is because users usually have diverse interests, and it is difficult for baselines to comprehensively model user interests. Different from these methods, *Uni-FedRec* learns multiple interest representations for a user from clusters of clicked news, which can comprehensively model diverse user interests in different fields. Besides, as shown in Table 3, *Uni-FedRec* can better protect user privacy than baseline methods. This is because baseline methods build user interest representations from the aggregation user’s clicked news, making user’s clicked news can be easily inferred from interest representations, which raises privacy leakage risks. Different from these methods, we propose to synthesize interest representations by combing privacy-insensitive basic interest embeddings shared among different users instead of user’s clicked news, which can better protect user privacy encoded in user representations.

4.3 News Recommendation Performance

Next, we combine different recall models and ranking models to evaluate the overall recommendation performance. We first use the recall model to generate 400 candidate news and further use the ranking model to rank these candidate news. We use users’ real click behaviors as ground truth and report AUC scores. Experimental results are presented in Table 4, and we only show results on *MIND* dataset

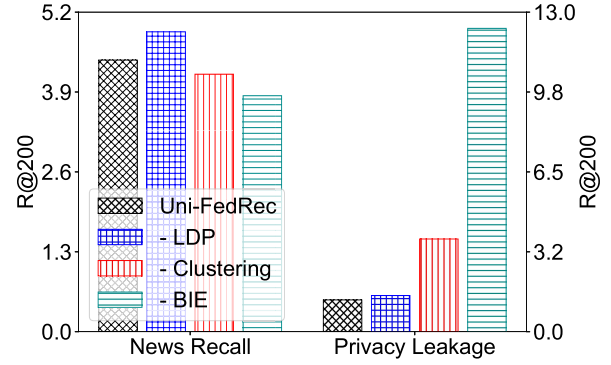


Figure 4: Ablation study on *Uni-FedRec*.

in the following sections due to space limitation. Results show that *Uni-FedRec* can consistently outperform baseline recall models when they are combined with a same ranking model. These results further verify that *Uni-FedRec* can outperform baseline methods in recommendation accuracy.

4.4 Ablation Study

As shown in Fig. 4, we verify the effectiveness of important modules in *Uni-FedRec* by removing them. First, after removing the LDP module in the recall model, news recall performance of *Uni-FedRec* improves while the privacy protection performance declines. This is intuitive since perturbation noise will make *Uni-FedRec* less accurate. Second, removing the hierarchical clustering framework hurts the news recall performance. This is because a user usually has diverse interests, which can be more comprehensively modeled by multiple interest representations. Third, removing *BIE* seriously hurts the privacy protection performance. This is because protected interest representations are synthesized from basic interest embeddings shared among different users, which contain much less private information of a specific user.

4.5 Influence of the LDP Noise

As shown in Fig. 5, we evaluate the influence of intensity λ_I of LDP noise n_I on *Uni-FedRec*. We find that with the increase of λ_I , news recall performance of *Uni-FedRec* declines and the privacy protection ability of *Uni-FedRec* increases. This is intuitive since incorporating larger noise will more seriously hurt the information capability of interest representations on both user interests and user privacy. Results in Fig. 5 inspire us that we can find a trade-off between recommendation accuracy and privacy protection by adjusting the intensity λ_I of LDP noise n_I on interest representations.




Figure 5: Influence of λ_I on *Uni-FedRec*.




Figure 6: Influence of d_c on *Uni-FedRec*.

4.6 Influence of Clustering Distance

In Fig. 6, we show the influence of clustering distance threshold d_c on *Uni-FedRec*. First, after d_c increases, recall performance of *Uni-FedRec* first increases. This is because small d_c makes *Uni-FedRec* build too many interest clusters, which may bring some noise and hurt the accuracy of interest representations. Second, when d_c becomes large enough, recall performance begins to decline. This is because larger d_c makes *Uni-FedRec* build fewer interest clusters and make it harder to comprehensively cover diverse user interests. Third, with the increase of d_c , the privacy protection performance of *Uni-FedRec* declines. This may be because when *Uni-FedRec* contains more interest channels, a single interest representation contains less private information. It may be easier for our proposed interest decomposer-aggregator method to protect private information encoded in them. Thus, a moderate value of d_c , i.e., 1, is suitable for *Uni-FedRec*.

4.7 Training Convergence Analysis

Fig. 7 shows the model training convergence of our recall model. The model is trained with different ratios of clients for a single federated model updating




Figure 7: Convergence curves in model training.

round. First, training of *Uni-FedRec* can usually converge in no more than two hundred steps, which verifies the efficiency of federated training of our recall model. Second, training convergence of *Uni-FedRec* will get faster and more stabilized if more clients can participate a single updating round. This is intuitive since updating parameters with more training data is usually more accurate.

5 Conclusion

In this paper, we propose a unified privacy-preserving news recommendation framework (*Uni-FedRec*) that can utilize user data locally stored in user clients to train models and serve users. Our framework contains a recall stage and a ranking stage. In the recall stage, the user client first employs a recall model to locally learn multiple interest representations from clicked news to model diverse user interest, which are further uploaded to the server to recall candidate news from a news pool. In the ranking stage, candidate news are distributed to the user client and locally ranked for personalized display. Besides, we propose an interest decomposer-aggregator method with permutation noise to protect private user information encoded in interest representations. In addition, *Uni-FedRec* collaboratively trains recall and ranking models on user data decentralized in massive user clients in a privacy-preserving way. Experiments on two real-world datasets show that our method can significantly outperform baseline methods and effectively protect user privacy.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant numbers U1936216, U1705261, and Tsinghua-Toyota Research Funds 20213930033.

References

- Mingxiao An, Fangzhao Wu, Chuhan Wu, Kun Zhang, Zheng Liu, and Xing Xie. 2019. Neural news recommendation with long-and short-term user representations. In *ACL*, pages 336–345.
- Sunil Arya, David M Mount, Nathan S Netanyahu, Ruth Silverman, and Angela Y Wu. 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *JACM*, pages 891–923.
- Trapit Bansal, Mrinal Das, and Chiranjib Bhattacharyya. 2015. Content driven user profiling for comment-worthy recommendations of news and blog articles. In *RecSys.*, pages 195–202.
- Di Chai, Leye Wang, Kai Chen, and Qiang Yang. 2019. Secure federated matrix factorization. *arXiv preprint arXiv:1906.05108*.
- Woo-Seok Choi, Matthew Tomei, Jose Rodrigo Sanchez Vicarte, Pavan Kumar Hanumolu, and Rakesh Kumar. 2018. Guaranteeing local differential privacy on ultra-low-power systems. In *ISCA*, pages 561–574.
- Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *RecSys.*, pages 191–198.
- Adrian Flanagan, Were Oyomno, Alexander Grigorievskiy, Kuan Eeik Tan, Suleiman A Khan, and Muhammad Ammad-Ud-Din. 2020. Federated multi-view matrix factorization for personalized recommendations. *arXiv preprint arXiv:2004.04256*.
- Suyu Ge, Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. Graph enhanced representation learning for news recommendation. In *WWW*, pages 2863–2869.
- Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178.
- Stephen C Johnson. 1967. Hierarchical clustering schemes. *Psychometrika*, pages 241–254.
- Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM*, pages 197–206.
- Dhruv Khattar, Vaibhav Kumar, Vasudeva Varma, and Manish Gupta. 2018. Weave&rec: A word embedding based 3-d convolutional network for news recommendation. In *CIKM*, pages 1855–1858.
- Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. 1997. GroupLens: applying collaborative filtering to usenet news. *Communications of the ACM*, pages 77–87.
- Guanyu Lin, Feng Liang, Weike Pan, and Zhong Ming. 2020. Fedrec: Federated recommendation with explicit feedback. *IEEE Intelligent Systems*.
- Jiahui Liu, Peter Dolan, and Elin Rønby Pedersen. 2010. Personalized news recommendation based on click behavior. In *IUI*, pages 31–40.
- Zheng Liu, Jianxun Lian, Junhan Yang, Defu Lian, and Xing Xie. 2020. Octopus: Comprehensive and elastic user representation for the generation of recommendation candidates. In *SIGIR*, pages 289–298.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pages 1273–1282.
- Khalil Muhammad, Qinqin Wang, Diarmuid O’Reilly-Morgan, Elias Tragos, Barry Smyth, Neil Hurley, James Geraci, and Aonghus Lawlor. 2020. Fedfast: Going beyond average for faster training of federated recommender systems. In *KDD*, pages 1234–1242.
- Shumpei Okura, Yukihiko Tagami, Shingo Ono, and Akira Tajima. 2017. Embedding-based news recommendation for millions of users. In *KDD*, pages 1933–1942.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Aditya Pal, Chantat Eksombatchai, Yitong Zhou, Bo Zhao, Charles Rosenberg, and Jure Leskovec. 2020. Pinnersage: Multi-modal user embedding framework for recommendations at pinterest. In *KDD*, pages 2311–2320.
- Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. 2021a. Personalized news recommendation with knowledge-aware interactive matching. In *SIGIR*, pages 61–70.
- Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. 2021b. Pp-rec: News recommendation with personalized user interest and time-aware news popularity. In *ACL*, pages 5457–5467.
- Tao Qi, Fangzhao Wu, Chuhan Wu, Yongfeng Huang, and Xing Xie. 2020. Privacy-preserving news recommendation model learning. In *EMNLP: Findings*, pages 1423–1432.
- Tao Qi, Fangzhao Wu, Chuhan Wu, Peiru Yang, Yang Yu, Xing Xie, and Yongfeng Huang. 2021c. Hierec: Hierarchical user interest modeling for personalized news recommendation. In *ACL*, pages 5446–5456.
- Xuebin Ren, Chia-Mu Yu, Weiren Yu, Shusen Yang, Xinyu Yang, Julie A McCann, and S Yu Philip. 2018. High-dimensional crowdsourced data publication with local differential privacy. *TIFS*, pages 2151–2166.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 6000–6010.

- Chong Wang and David M Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *KDD*, pages 448–456.
- Heyuan Wang, Fangzhao Wu, Zheng Liu, and Xing Xie. 2020. Fine-grained interest matching for neural news recommendation. In *ACL*, pages 836–845.
- Qinyong Wang, Hongzhi Yin, Tong Chen, Junliang Yu, Alexander Zhou, and Xiangliang Zhang. 2021. Fast-adapting and privacy-preserving federated recommender system. *arXiv preprint arXiv:2104.00919*.
- Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019a. Neural news recommendation with attentive multi-view learning. *IJCAI*, pages 3863–3869.
- Chuhan Wu, Fangzhao Wu, Mingxiao An, Yongfeng Huang, and Xing Xie. 2019b. Neural news recommendation with topic-aware news representation. In *ACL*, pages 1154–1159.
- Chuhan Wu, Fangzhao Wu, Mingxiao An, Tao Qi, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019c. Neural news recommendation with heterogeneous user behavior. In *EMNLP*, pages 4876–4885.
- Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. 2021a. Fedgmn: Federated graph neural network for privacy-preserving recommendation. *arXiv preprint arXiv:2102.04925*.
- Chuhan Wu, Fangzhao Wu, Tao Di, Yongfeng Huang, and Xing Xie. 2020a. Fedctr: Federated native ad ctr prediction with multi-platform user behavior data. *arXiv preprint arXiv:2007.12135*.
- Chuhan Wu, Fangzhao Wu, Suyu Ge, Tao Qi, Yongfeng Huang, and Xing Xie. 2019d. Neural news recommendation with multi-head self-attention. In *EMNLP*, pages 6390–6395.
- Chuhan Wu, Fangzhao Wu, and Yongfeng Huang. 2021b. Personalized news recommendation: A survey. *arXiv preprint arXiv:2106.08934*.
- Chuhan Wu, Fangzhao Wu, Yongfeng Huang, and Xing Xie. 2021c. User-as-graph: User modeling with heterogeneous graph pooling for news recommendation. In *IJCAI*, pages 1624–1630.
- Chuhan Wu, Fangzhao Wu, Junxin Liu, and Yongfeng Huang. 2019e. Hierarchical user and item representation with three-tier attention for recommendation. In *NAACL*, pages 1818–1826.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020b. Sentirec: Sentiment diversity-aware neural news recommendation. In *AACL*, pages 44–53.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020c. User modeling with click preference and reading satisfaction for news recommendation. In *IJCAI*, pages 3023–3029.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2021d. Empowering news recommendation with pre-trained language models. In *SIGIR*, pages 1652–1656.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2021e. Two birds with one stone: Unified model learning for both recall and ranking in news recommendation. *arXiv preprint arXiv:2104.07404*.
- Chuhan Wu, Fangzhao Wu, Tao Qi, Jianxun Lian, Yongfeng Huang, and Xing Xie. 2020d. Ptum: Pre-training user model from unlabeled user behaviors via self-supervision. In *EMNLP: Findings*, pages 1939–1944.
- Chuhan Wu, Fangzhao Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2021f. Fairrec: fairness-aware news recommendation with decomposed adversarial learning. In *AAAI*, pages 4462–4469.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020e. Mind: A large-scale dataset for news recommendation. In *ACL*, pages 3597–3606.
- Enyue Yang, Yunfeng Huang, Feng Liang, Weike Pan, and Zhong Ming. 2021. Fcmf: Federated collective matrix factorization for heterogeneous collaborative filtering. *KBS*, page 106946.
- Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. In *NeurIPS*, pages 4876–4885.