



 Latest updates: <https://dl.acm.org/doi/10.1145/3637528.3671984>

RESEARCH-ARTICLE

A Population-to-individual Tuning Framework for Adapting Pretrained LM to On-device User Intent Prediction

JIAHUI GONG, Beijing National Research Center for Information Science and Technology, Beijing, China

JINGTAO DING, Beijing National Research Center for Information Science and Technology, Beijing, China

FANJIN MENG, Beijing National Research Center for Information Science and Technology, Beijing, China

GUILONG CHEN, Honor Device Co., Ltd., Shenzhen, Futian, China

HONG CHEN, Honor Device Co., Ltd., Shenzhen, Futian, China

SHEN ZHAO, Honor Device Co., Ltd., Shenzhen, Futian, China

[View all](#)

Open Access Support provided by:

Beijing National Research Center for Information Science and Technology

Honor Device Co., Ltd.



PDF Download
3637528.3671984.pdf
18 January 2026
Total Citations: 7
Total Downloads: 885

Published: 25 August 2024

[Citation in BibTeX format](#)

KDD '24: The 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining

August 25 - 29, 2024
Barcelona, Spain

Conference Sponsors:

SIGMOD
SIGKDD

A Population-to-individual Tuning Framework for Adapting Pretrained LM to On-device User Intent Prediction

Jiahui Gong

Department of Electronic
Engineering, BNRist, Tsinghua
University
Beijing, China
gjh22@mails.tsinghua.edu.cn

Guilong Chen

Honor Device Co., Ltd.
Shenzhen, China
chenguilong@honor.com

Haisheng Lu

Honor Device Co., Ltd.
Shenzhen, China
luhaisheng@honor.com

Jingtao Ding*

Department of Electronic
Engineering, BNRist, Tsinghua
University
Beijing, China
dingjt15@tsinghua.org.cn

Hong Chen

Honor Device Co., Ltd.
Shenzhen, China
chenhong3@honor.com

Fanjin Meng

Department of Electronic
Engineering, BNRist, Tsinghua
University
Beijing, China
mengfj23@mails.tsinghua.edu.cn

Shen Zhao

Honor Device Co., Ltd.
Shenzhen, China
zhaoshen@honor.com

Yong Li

Department of Electronic
Engineering, BNRist, Tsinghua
University
Beijing, China
liyong07@tsinghua.edu.cn

ABSTRACT

Mobile devices, especially smartphones, can support rich functions and have developed into indispensable tools in daily life. With the rise of generative AI services, smartphones can potentially transform into personalized assistants, anticipating user needs and scheduling services accordingly. Predicting user intents on smartphones, and reflecting anticipated activities based on past interactions and context, remains a pivotal step towards this vision. Existing research predominantly focuses on specific domains, neglecting the challenge of modeling diverse event sequences across dynamic contexts. Leveraging pre-trained language models (PLMs) offers a promising avenue, yet adapting PLMs to on-device user intent prediction presents significant challenges. To address these challenges, we propose PITuning, a Population-to-Individual Tuning framework. PITuning enhances common pattern extraction through dynamic event-to-intent transition modeling and addresses long-tailed preferences via adaptive unlearning strategies. Experimental results on real-world datasets demonstrate PITuning's superior intent prediction performance, highlighting its ability to capture long-tailed preferences and its practicality for on-device prediction scenarios.

CCS CONCEPTS

- Information systems → Recommender systems; Personalization;
- Computing methodologies → Machine learning.

* Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

KEYWORDS

Device-cloud collaboration; Pretrained language model; Personalization; User intent

ACM Reference Format:

Jiahui Gong, Jingtao Ding*, Fanjin Meng, Guilong Chen, Hong Chen, Shen Zhao, Haisheng Lu, and Yong Li. 2024. A Population-to-individual Tuning Framework for Adapting Pretrained LM to On-device User Intent Prediction. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), August 25–29, 2024, Barcelona, Spain*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3671984>

1 INTRODUCTION

Nowadays mobile devices, especially smartphones, have become a major object that individuals interact with in their daily lives. For example, users use their phones to monitor sleep, wake themselves up, hail a car for commuting, watch short videos in rest time, pay money at restaurants, etc., across most activities in one day. Empowered by the recent booming of generative artificial intelligence (AI) services (e.g., chatGPT [3]), the smartphone can further evolve into a personalized assistant that can perceive user needs in advance and timely schedule corresponding services. The key pathway toward this future is the capability to predict smartphone users' intents, which refers to what activity they intend to do, based on their previous action sequences and contextual information [16, 17].

Existing works mostly focus on predicting user intents within one specific domain, for example, purchase intent in online platforms [15, 18, 26, 43], search intent in search engines [37, 48], pedestrian intention for robots or autonomous vehicles [1, 34]. To characterize complex dependencies between intent and context, they leverage specific network architectures including feature interaction networks [27, 41] or graph neural networks [15, 43]. In

contrast, predicting users' daily activity intent when using smartphones requires modeling diverse event sequences across dynamic changing contexts, which generally rely on large-scale behavioral data. However, with increasing concerns about data privacy leakage and real-time serving latency, real-world prediction applications usually adopt on-device model training and deployment, which adds constraints on data scales and exacerbates the data lacking issue.

Pretrained language models (PLMs) [3, 29], on the other hand, provide a promising solution owing to their encoded knowledge and commonsense reasoning capability acquired through extensive training on diverse datasets. For example, if someone talks about going jogging every morning, a PLM can infer that the individual values fitness, which might predict other health-related behaviors. In this regard, PLMs have been successfully adapted to other cross-domain tasks related to human behaviors, like recommender systems [2, 7, 45] and mobility trajectories [11, 20, 38]. Therefore, we propose to leverage PLMs for on-device user intent prediction, i.e., adapting PLMs from the language domain into the daily human behavior domain, which is non-trivial due to the following three challenges:

- **Population-level common behavioral patterns are hard to extract from the noisy aggregation of diverse event sequences.** Predicting user intent based on their previous action events requires the characterization of common transition patterns from event sequences to specific intent. However, not all events correlate to the generation of intent, i.e., information redundancy, and this changes with intent type. Although transformer-based architecture has proven its usefulness in the sequential modeling of user events [30, 36], it remains questionable whether common event-intent transition patterns shared among the population can be extracted from the above noisy and redundant event sequences.
- **Individual-level long-tailed preferences are hard to capture by large LMs.** Besides common behavioral patterns, individual preference also matters a lot in predicting user intent. For example, compared with public transport, car-hailing might be a long-tailed choice globally, while favored by a few users. However, long-tailed individual preferences are prone to be overtaken by population-level patterns that dominate the population's behavioral data. This inevitably leads to a biased model favoring those intents with a high proportion after tuning a PLM. Existing works on aligning LM for behavioral modeling [35] tend to construct tuning tasks analogous to their counterparts in NLP like prompt tuning [7] or instruction tuning [2, 53]. Without a specific design, however, it is generally difficult to alleviate the above bias problem of the long-tailed preferences given rather limited individual behavioral data.
- **Designing a practical LM tuning framework to support on-device learning and inference of user intent is difficult.** Existing works have proposed a few cloud-device collaboration approaches to achieve on-device prediction or recommendation, mainly targeting device-side personalization. Differently, the expected tuning framework is tailored for pretrained LMs and should be able to leverage large-scale

population-level data efficiently and limited individual-level data effectively.

In this paper, we propose a novel **Population-to-Individual Tuning** framework (named PITuning) for adapting pretrained LM to on-device user intent prediction. The core of the PITuning framework is the population-level behavioral data tuning on a pretrained LM that produces a powerful but gigantic global predictor at the cloud side, and the individual-level tuning that adaptively distills this predictor into a lightweight user-specific predictor at the device side. To solve the first challenge of extracting common behavioral patterns at the population level, PITuning is designed to better capture dynamic event-to-intent transition patterns, i.e., event-wise information enhancement by an auxiliary event-reconstruction loss and intent-wise attentive modeling on top of pretrained transformer. As for the second challenge of capturing long-tailed preference distribution at the individual level, PITuning is equipped with a novel unlearning strategy for each user that first identifies a set of intents, which are under-represented in population data but emphasizes unique preferences of this user and then remove the model's memorization on these intents. This further guarantees effectively capturing long-tailed preference by tuning on individual behavioral data at the device side. To summarize, our main contributions are as follows.

- We provide a novel angle of adapting PLMs into the human behavioral domain and further resolve the longstanding issue of capturing long-tailed user preferences.
- We design a population-to-individual tuning framework for PLM that extracts common behavioral patterns and captures individual unique patterns simultaneously, compatible with on-device prediction scenarios.
- Experiment results on two real-world datasets demonstrate the superiority of our PITuning over state-of-the-art baselines in terms of intent prediction performance. Notably, the outperformance regarding the macroscopic average of precision and recall is 24%-37%, underscoring its capability of capturing long-tailed preference for individuals. Ablation studies and in-depth analysis further support the rationality behind specific method design, as well as the high practicality in terms of efficiency and scalability.

2 PRELIMINARY

2.1 Data Analysis

We begin with a comprehensive data analysis. Initially, we randomly sample 1,000 users to calculate their intent distribution. Subsequently, we employ the KMeans method [9] to cluster users' intent distribution and visualize the result using t-SNE [39], as illustrated in Figure 1. Additionally, we present the population-level intent distribution alongside the distribution for each cluster. From the figure, we observe that intent distribution varies significantly between clusters. This discrepancy undoubtedly complicates the task of user-personalized modeling.

2.2 Problem Statement

Now we give a formal definition of our research problem:

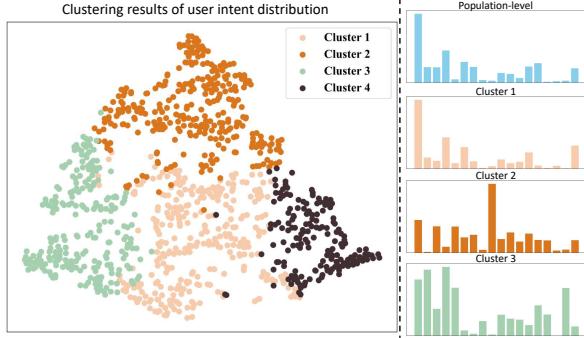


Figure 1: Distribution gap exists between population-level preference and individual-level preference (by comparing frequency histogram of user intent).

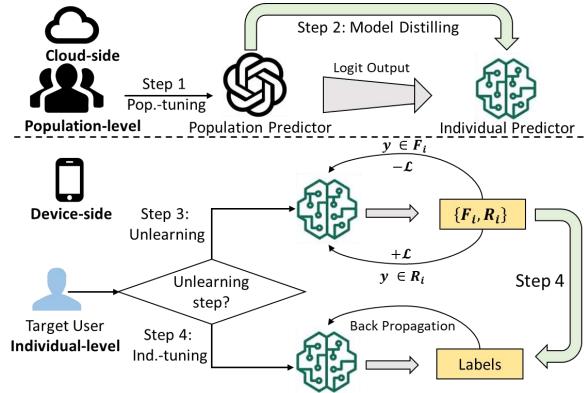


Figure 2: The workflow of PITuning framework.

PROBLEM1 (User intent prediction). The behavior corresponding to the i -th user intention can be represented as $x_i = (u_i, l_i, t_i, e_i)$, indicating that a specific event e_i takes place involving user u_i at location l_i during time slot t_i . Here, u_i , l_i , t_i , and e_i refer to the user ID, location ID, time slot ID, and event ID, respectively. We use $\mathcal{U}, \mathcal{L}, \mathcal{T}, \mathcal{E}$ to denote the sets of users, locations, time slots, and events, with their respective sizes given by N_U, N_L, N_T , and N_E . As outlined in the introduction, each user exhibits a particular intention y_i associated with an event-related behavior x_i . We define \mathcal{I} as the set of possible intentions, with its size represented by N_I . The event encompasses specific instances involving users, such as the use of app services, spatial trajectory occurrences, and system-related events. The intent captures the underlying goal, purpose, or objective driving these events, effectively grouping them into categories. Therefore, the quantity of distinct intents, denoted by N_I , is typically less than the total count of events, represented by N_E .

User intent prediction aims to forecast future user intent based on its past I event series, which can be formed as,

$$y_t = f(x_{t-I}, x_{t-I+1}, \dots, x_{t-1}) \quad (1)$$

3 METHOD

3.1 Framework Overview

We introduce our PITuning framework for adapting PLM to on-device intent prediction, as depicted in Figure 2. On the cloud side,

we utilize aggregated behavioral data collected from a population to fine-tune a global predictor, capturing population-level common behavioral patterns. Subsequently, we perform model distillation to obtain a lightweight predictor suitable for on-device deployment. On the device side, before further fine-tuning on individual data, we incorporate a novel unlearning strategy to identify and mitigate biases resulting from uneven learning of intents during population-level tuning. Finally, after two stages of PITuning, we attain a lightweight yet personalized model capable of accurate and efficient intent prediction on the device.

3.2 Population-level Tuning

We leverage population data alongside a PLM to model common behavioral patterns. The architecture of our model is depicted in Figure 3(a), where we integrate parameters from GPT2 [29], an NLP pre-trained transformer model. Additionally, to enhance learning of event-to-intent transition patterns shared among the population, we introduce a masked event-reconstruction loss at the event level and utilize intent-wise attentive modeling atop the pretrained transformer. Finally, we distill a lightweight predictor under the guidance of the global predictor to meet deployment requirements on the device.

3.2.1 Embedding layer. Since we apply the NLP pre-trained model to a new modality, we create four embedding layers to get the location embedding $E_l \in \mathbb{R}^{I \times d}$, weekday embedding $E_w \in \mathbb{R}^{I \times d}$, time-slot embedding $E_t \in \mathbb{R}^{I \times d}$ and event embedding $E_e \in \mathbb{R}^{I \times d}$ respectively, where d denotes the embedding size.

3.2.2 Transformer block. In the global predictor, we employ the GPT2 model as the foundation for our transformer blocks. Trained on various web data, the GPT2 model is imbued with extensive knowledge, common sense, and fundamental principles, showing a robust capacity for generalization. We concat the location embedding, weekday embedding, time-slot embedding, and event embedding and put them into transformer blocks to obtain an implicit representation of the historical event sequence $H_t \in \mathbb{R}^{I \times 4d}$, which can be formed as,

$$H_t = \text{GPT2}(\text{concat}(E_l, E_w, E_t, E_e)). \quad (2)$$

3.2.3 Intent-aware attentive modeling. Notice that different intents exhibit preferences for varying lengths of historical data. To address this, we have developed a novel Intention Attention Network (IAT) that introduces a novel designed local activation unit to adaptively weigh sequences of historical events, accommodating the unique requirements of each intent. We create the learnable intent embedding matrix $E_i \in \mathbb{R}^{N_I \times 4d}$ and feed it into IAT together with the history matrix.

Specifically, we apply activation units to the features derived from users' historical behaviors. These units function through a weighted sum pooling mechanism to adaptively compute the representation of intents, as detailed in Equation 3,

$$\mathbf{H}_w = \text{IAT}(\mathbf{E}_i, \mathbf{H}_t) = \sum_{j=0}^I a(h_j, E_i) h_j = \sum_{j=0}^I w_j h_j. \quad (3)$$

Through this approach, \mathbf{H}_w changes across different intents, where $a(\cdot)$ represents a feed-forward network that yields activation weights.

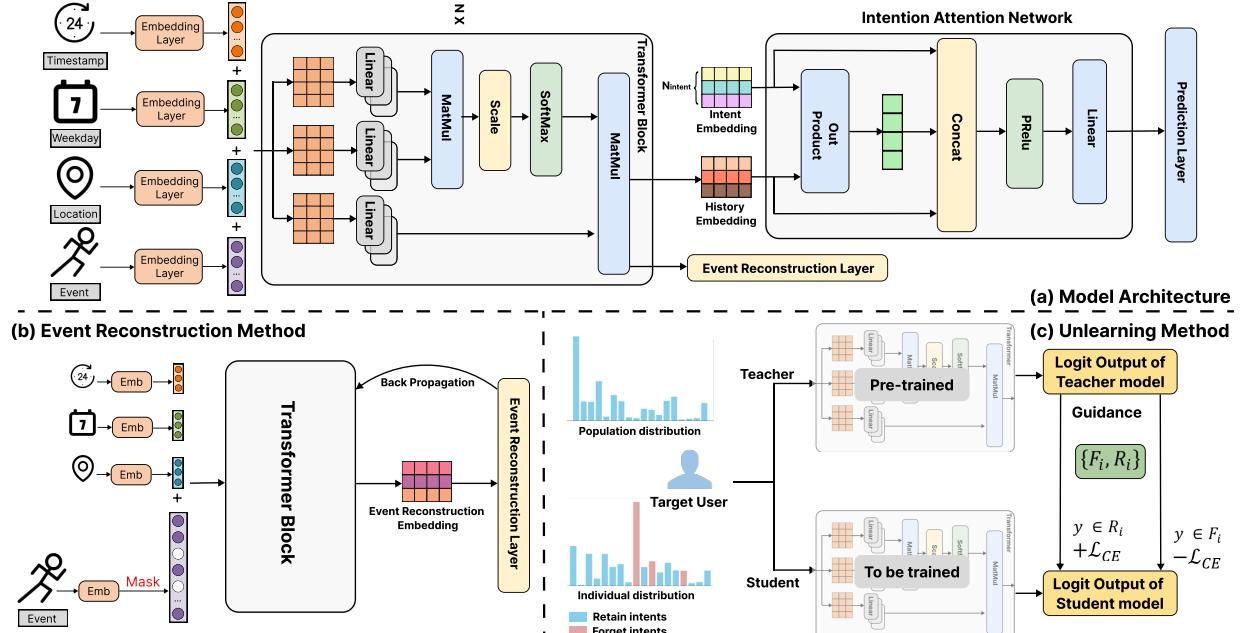


Figure 3: (a) The intent predictor architecture. (b) The masked event reconstruction in the population-level tuning. (c) The adaptive unlearning in the individual-level tuning.

These weights are then combined through an outer product operation and integrated into the subsequent network layers to enhance relevance modeling.

Next, we use a Multilayer Perception (MLP) to be the prediction layer, which can be formed as,

$$\mathbf{m} = f(\mathbf{H}_w) = \mathbf{W}_2(\sigma(\mathbf{W}_1 \mathbf{H}_w + b_1)) + b_2, \quad (4)$$

where \mathbf{W}, b are the trainable weight matrix and the bias matrix. The output of the MLP is the predicted intent distribution.

3.2.4 Event-reconstruction auxiliary loss. To improve the model's proficiency in accurately capturing event-to-intent transition patterns, we employ a masked event reconstruction loss [33], which reconstructs the original event sequences based on the given partially observed signals, as shown in Figure 3(b). Specifically, we randomly mask the event embedding E_m , and input them into the GPT2 model according to 2. Next, we employ an MLP to be the event reconstruction layer to reconstruct the event sequence. The cross-entropy loss function is then used to assist model training. The loss function in population-level tuning can be formed as,

$$\mathcal{L}_{pop} = \mathcal{L}_M + \mathcal{L}_{CE}(\mathbf{m}, L) = \mathcal{L}_{CE}(e_i, e_m) + \mathcal{L}_{CE}(\mathbf{m}, L), \quad (5)$$

where e_i denotes the original event sequence, e_m denotes the predicted event sequence, and L denotes the ground truth of input.

3.2.5 Model distilling. To meet the requirement of deployment, we utilize the model distilling method, which is to train a smaller model (called the student model) to imitate the behavior of a larger model (called the teacher model). The details of the model distillation process are shown in Appendix A.

To guide the training of the student model, we design the soft loss for the soft targets, which is the Kullback-Leibler Divergence between the logit output of the teacher and the student network.

Meanwhile, we also utilize the cross-entropy loss to ensure the student model learns the correct classifications. The loss function can be formed as follows,

$$\mathcal{L}_{soft} = \text{KL}\left(\frac{\mathbf{M}_s}{\epsilon}, \frac{\mathbf{M}_t}{\epsilon}\right), \quad \mathcal{L}_{hard} = \mathcal{L}_{CE}(\mathbf{M}_s, L) \quad (6)$$

$$\mathcal{L}_D = \alpha \mathcal{L}_{soft} + (1 - \alpha) \mathcal{L}_{hard} \quad (7)$$

where $\mathbf{M}_t, \mathbf{M}_s$ denotes the logits output of the teacher and student model respectively, and ϵ is a hyper-parameter, which means the temperature to smooth the probability distribution, while α is a hyper-parameter to balance the importance of two loss functions. By doing so, the student model learns both the fine-grained information from the teacher model's output and the essential classification ability, resulting in a smaller, more efficient model that retains much of the teacher model's predictive power. Subsequently, the student model is deployed on the device side.

3.3 Individual-level Tuning

In the individual-level tuning, we harness personalized user data to fine-tune the model, enabling it to adapt to and learn individual preferences. However, the key challenge is that there may be a significant difference between population intent distribution and individual intent distribution (as shown in our previous data analysis in Figure 1), leading to a bias, particularly for some long-tail intents during the population tuning stage. Therefore we first design an adaptive unlearning strategy to help the model disregard these biases. After that, we can finetune a personalized model that is both accurate and efficient.

3.3.1 Adaptive unlearning on biased intents. Unlearning involves intentionally disregarding or ignoring specific data or patterns in a trained neural network [4]. Initially, we decide whether each user's

Algorithm 1 Population-to-Individual Tuning Framework**Population-level tuning****Require:** Population data $x_i = (l_i, w_i, t_i, e_i)$ **Ensure:** The lightweight predictor M_l

- 1: $\mathbf{E}_l, \mathbf{E}_w, \mathbf{E}_t, \mathbf{E}_e \leftarrow emb(l_i), emb(w_i), emb(t_i), emb(e_i)$ \triangleright *Input Embedding.*
- 2: $\mathbf{H}_t = GPT2(\text{concat}(\mathbf{E}_l, \mathbf{E}_w, \mathbf{E}_t, \mathbf{E}_e))$
- 3: $\mathbf{M} = \text{MLP}(\text{IAT}(\mathbf{H}_t))$ \triangleright *Population Intent Prediction.*
- 4: $\mathbf{E}_m = \text{Mask}(\mathbf{E}_e)$ \triangleright *Mask the event sequence.*
- 5: $\mathbf{H}_m = GPT2(\text{concat}(\mathbf{E}_l, \mathbf{E}_w, \mathbf{E}_t, \mathbf{E}_m))$
- 6: $\mathbf{M}_m = \text{MLP}(\mathbf{H}_t)$ \triangleright *Event Reconstruction.*
- 7: Model Distillation to obtain the lightweight predictor M_l .

Individual-level tuning**Require:** Individual data $x_j = (l_j, w_j, t_j, e_j)$ of User j **Ensure:** The tuned lightweight predictor M_f

- 1: $(F_i, R_i) \leftarrow \text{ManageIntents}(P_{pop}, P_{in})$
- 2: **if** $F_j \neq \emptyset$ **then**
- 3: the adaptive unlearning to forget F_j
- 4: **end if**
- 5: Finetune the lightweight predictor M_f

intent should be forgotten or retained. We propose two methods to identify the forgotten intents. First, we analyze the intent distribution of the global predictor output at the population level. If the proportion of an intent $P_{out}(i)$ is less than the threshold ε , it is part of the static forgotten set F_{sta} ,

$$I_i \in F_{sta} \text{ if } P_{out}(I_i) < \varepsilon, \quad (8)$$

where I_i denotes the intents i . Secondly, if the proportion of intent at the population level P_{pop} is less than the average while the proportion at the individual level P_{in} is greater than the average, then it belongs to the dynamic forgotten set F_{dyn} ,

$$I_i \in F_{dyn} \text{ if } P_{pop} < \frac{1}{N_I} \text{ and } P_{in} > \frac{1}{N_I}. \quad (9)$$

Finally, the union of static forgotten set and dynamic forgotten set is taken as the forgotten set, which can be formed as,

$$F_i = F_{dyn} \cup F_{sta}, R_i = \text{others}, \quad (10)$$

where F_i, R_i represents the forgotten set and retained set of user i . If it is determined that intent needs to be forgotten by the user, the adaptive unlearning is applied; if not, it is deemed unnecessary. To effectively achieve the unlearning goal, we design an unlearning loss:

$$\mathcal{L}_{un} = \lambda \mathcal{L}_{CE}(\mathbf{M}_r, R_i) - \mathcal{L}_{CE}(\mathbf{M}_f, F_i), \quad (11)$$

where \mathbf{M}_f signifies the model's output for the intent category designated for forgetting, and λ is a hyper-parameter to balance the trade-off between forgetting and retaining. Intuitively, during the unlearning process, the model is learned to minimize the loss between the output from the updated model and the original model on the intent to retain while maximizing the loss between the output from them on the data to forget.

Table 1: Statistics of the datasets used in our experiments.

Datasets		Honor Dataset	Mobile Dataset
	Type of Events	114	12
	Type of Intents	18	12
Population level	Users	4,500	4,000
	Duration	6.1-8.22, 2023	10.1-31, 2016
	Number of logs	10,376,148	334,651
Individual level	Users	5,000	2,000
	Duration	8.23-9.10, 2023	10.1-31, 2016
	Number of logs	976,788	208,161

3.3.2 *Finetuning for personalized model.* Finally, we utilize personalized individual data to fine-tune the model for each user. This fine-tuning process enables the model to transition from capturing common behavioral patterns to reflecting a user's unique preferences, thereby enhancing the accuracy of the model's predictions. We also use cross-entropy loss to guide model tuning.

4 EXPERIMENT**4.1 Experiment Settings**

4.1.1 *Datasets.* We evaluate the performance of our model on two large-scale real-world activity datasets.

- **Honor Dataset.** The Honor Dataset is sampled from the usage log of the mobile phones. When a user uses mobile phones, various types of logs are generated, desensitized and reported (with user consent). We selected 114 types of events that are commonly monitored in most mobile applications and classified them into 18 intents, which cover the aspects of news, study, work, entertainment, sports, etc. We sampled two datasets between June 1st and August 22nd, 2023 (the first) and August 22nd and September 10th, 2023 (the second) which in total contain 4,500 and 5,000 anonymous users.

- **Mobile Dataset.** The Mobile Dataset consists of anonymous user trajectory data collected by a major mobile network operator in China in October. The dataset comprises 6,000 users, of which, at the population level, we select 4,000 users for training, and at the individual level, we select the remaining users. In this dataset, we use the location category as the activity and intent type.

Table 1 shows the statistics of the Honor dataset and Mobile dataset. The large-scale and fine-grained datasets can ensure the validity of the model test.

4.1.2 *Metrics.* To assess model performance, we employ five widely used metrics: weighted precision ($Prec_w$), weighted recall (Rec_w), macro precision ($Prec_m$), macro recall (Rec_m), and NDCG(N). Weighted metrics and NDCG gauge classification accuracy and ranking quality, respectively, while macro metrics evaluate the average prediction accuracy for each intent, indicating the model's predictive quality across intents. A smaller gap between weighted and macro metrics implies consistent prediction accuracy across intents, reflecting fairness. Conversely, a large gap suggests inadequate modeling of long-tail intents, leading to suboptimal outcomes. Refer to Appendix D for metric calculations.

Table 2: Overall prediction performance PITuning compared with baselines on Honor and Mobile datasets.

Model	Honor Dataset						Mobile Dataset					
	$Prec_w$	Rec_w	$Prec_m$	Rec_m	N@3	N@5	$Prec_w$	Rec_w	$Prec_m$	Rec_m	N@3	N@5
CLOVER	0.4479	0.4516	0.2494	0.2527	0.6094	0.6213	0.7052	0.7505	0.6004	0.6350	0.7860	0.8270
MetaBert4Rec	0.4683	0.5028	0.2888	0.3218	0.6842	0.7023	0.7714	0.8174	0.6419	0.6739	0.8268	0.8814
P5	0.4722	0.5161	0.2452	0.2807	0.6045	0.6284	0.7436	0.7930	0.6133	0.6410	0.8104	0.8672
InstructRec	0.4315	0.4680	0.2423	0.2684	0.6407	0.6724	0.7318	0.7743	0.6043	0.6397	0.7968	0.8444
LSAT	0.4714	0.5008	0.2778	0.2983	0.6055	0.6301	0.7572	0.8058	0.6415	0.6845	0.8220	0.8809
OFA	0.4928	0.5243	0.3366	0.3756	0.7032	0.7244	0.7851	0.8258	0.6538	0.7062	0.8488	0.9024
TallRec	0.4486	0.4764	0.2659	0.3083	0.5972	0.6102	0.7200	0.7693	0.6261	0.6532	0.7949	0.8345
EODRec	0.4517	0.4958	0.2799	0.2903	0.5867	0.6039	0.7489	0.8087	0.6072	0.6519	0.8296	0.8939
MPDA	0.4947	0.5197	0.3408	0.3841	0.7117	0.7371	0.7785	0.8361	0.6581	0.7085	0.8542	0.9185
ours	0.5374	0.5599	0.4693	0.4840	0.7329	0.7626	0.8715	0.9002	0.8449	0.8802	0.9506	0.9537
Improv.	8.63%	6.79%	37.71%	26.01%	2.98%	3.46%	11.00%	7.67%	28.38%	24.23%	11.29%	3.83%

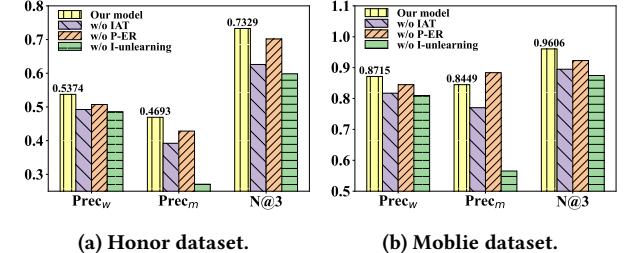
4.1.3 Baselines. We elaborately select the following nine representatives to be compared with our proposed algorithms, which cover the meta-learning methods for personalized recommendations (CLOVER [44], MetaBert4Rec [13]), LLM-based recommendations (P5 [7], InstructRec [53], LSAT [35], One fits All (OFA) [38], TallRec [2]) and device-cloud collaboration recommendations (EODRec [46], MPDA [47]). We provide the details of baselines in Appendix C.

4.1.4 Implementation Details. Our model employs the Adam optimizer with a learning rate of 0.01 across two tuning phases. And we set the input length as 30. During the population-level tuning stage, we utilize the GPT2 small version for the transformer blocks, which features a 12-layer transformer architecture and a 768-dimensional feature space, while on the individual-level tuning stage, we choose a transformer decoder with 4-layer and a 768-dimensional feature space to meet the deployment requirement. As for hyper-parameters α , ϵ , and λ , we set 0.5, 1, and 1 respectively. Details of hyperparameters are shown in Appendix B. In the honor dataset, one week is allocated for training, one day for validation, and four days for testing. In the mobile dataset, 60% of the data is used for training, with 10% for validation and 30% for testing. The code is available at <https://github.com/tsinghua-fib-lab/LLM-for-User-Intent>.

4.2 Overall Performance

In Table 2, we display the overall results of our model, meta-learning methods (CLOVER, MetaBert4Rec), LLM-based recommendations (P5, InstructRec, LSAT, OFA, TallRec) and device-cloud collaboration recommendations (EODRec, MPDA) to predict the next user intention in two datasets. We list three metrics of all methods. From the result, we have the following findings:

- Our framework steadily achieves the best performance.** Our model gets superior results on both datasets and performs better than other compared algorithms. For example, the macro metrics improvement of our model is around 24% to 37% compared with the second-best performance model (MPDA). The NDCG improvement of our model is about 3% to 11%.

**Figure 4: Ablation study.**

- Our model has the smallest difference between weighted metrics and macro metrics.** Weighted metrics and macro metrics count the global accuracy and the average accuracy of each intent respectively. A smaller difference suggests comparable prediction accuracy across different intents, indicating fairness among the intents. Our model utilizes adaptive unlearning to effectively correct the long-tail intent learning bias caused by the model in population-level tuning, and improve the accuracy.
- The method of LLM-based model with device-cloud collaboration is necessary for user behavior modeling.** MPDA utilizes the LLM-based model with the device-cloud collaboration method resulting in the best performance within the baseline. However, MPDA fails to fully exploit the benefits of diverse user data available on the cloud side. In contrast, our proposed PITuning framework captures generalized user behavior patterns, leading to superior performance.

4.3 Ablation Study

To gain a deeper understanding of each component of our model, we carried out a sequence of ablation studies. Firstly, we removed the Intention Attention Network (IAT) within this model, followed by removing the event reconstruction loss in the population-level tuning (P-ER). Subsequently, we removed the adaptive unlearning in the individual-level tuning(I-unlearning).

The results of the ablation study are presented in Figure 4. We observed that the absence of the Intention Attention Network (IAT)

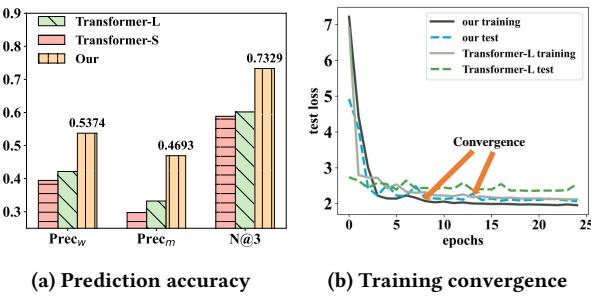


Figure 5: Comparing performance without pretrained LM.

hindered the model's ability to appropriately assign weights to each intent, consequently impacting both $Prec_w$ and $Prec_m$. Additionally, the event reconstruction loss played a pivotal role in guiding the transformer block towards more accurate modelling of users' historical event sequences, thereby enhancing the model's performance. Furthermore, we noticed that the omission of the adaptive unlearning compromised the model's capacity to effectively handle long-tail intents, resulting in a significant reduction in $Prec_m$ by approximately 37%.

4.4 Analysis of Population-level Tuning

- Performance improvement brought by pretrained LM.

We assessed how LLM contributes to modeling population-level common behavior patterns by replacing the LLM with transformer encoders of two sizes. One matches the size of GPT2 (Transformer-L), while the other matches the size of the lightweight predictor we distilled (Transformer-S), allowing them to train from scratch.

Figure 5 compares the prediction accuracy and loss among the models. Our analysis underscores that without leveraging the PTM, the model lacked foundational common sense and rule-based guidance, leading to a significant decline in its ability to capture common behavioral patterns and accuracy. Additionally, the Transformer-S, with fewer parameters, encountered challenges in modeling complex user behaviors, resulting in inferior performance. Moreover, with the guidance of LLM, the model demonstrated faster convergence.

- Effectiveness of extracting common behavioral pattern.

To highlight the effectiveness of IAT in capturing intent-aware transition patterns in user event sequences, we visualize the attention map between intents and historical sequences. Figure 6 presents the resulting attention maps, highlighting the IAT's ability to discern each intent's preference for historical sequence length. Analysis of the attention map reveals that certain intents, such as short video, game, and photo intents, predominantly rely on short-term historical sequences. Conversely, intents like checking the weather, taking a taxi, and exercising necessitate long-term historical sequences. Additionally, some intents rely on both short-term and long-term historical sequences, such as checking the weather, music, and audiobook intents. These insights uncover users' daily behavior patterns, enabling researchers to construct more nuanced historical sequences and features to enhance accuracy.

To showcase the effectiveness of the event reconstruction loss, we compare the difference between the intent distribution output by our model and OFA during the population-level tuning stage.

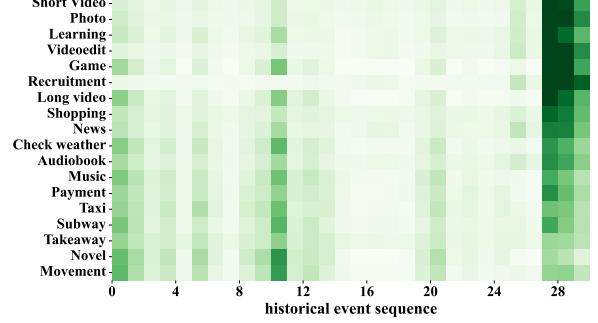


Figure 6: The attention map of different intents illustrating diverse event-to-intent transition patterns.

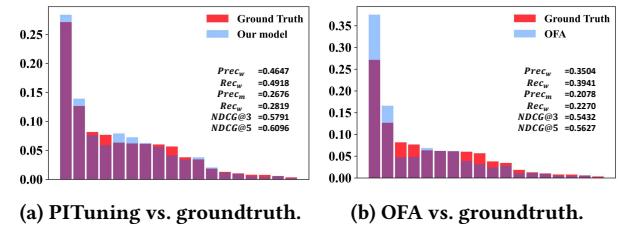


Figure 7: Comparison of intent distribution generated by PITuning and OFA after population-level tuning.

Figure 7 presents the results, demonstrating that our model could well model transitions from events to intents with the help of the event reconstruction loss. The closer the output intent distribution is to the real distribution, the more conducive it is to capture the common behavior patterns in the population-level tuning.

4.5 Analysis of Individual-level Tuning.

- Choice of device model.

To evaluate the efficiency of the lightweight predictor obtained through model distillation, we compared it with the tree model LightGBM [12], and two variants of the original population predictor at the cloud side, i.e., full-parameter tuned population predictor (FP), and partial parameter tuned population predictor (PP). Additional details about the tree model can be found in Appendix E. PP is inspired by the cross-domain adaptation techniques used in OFA [38], which argues that self-attention layers and feed-forward neural networks encapsulate most learned knowledge and can be frozen during the finetuning process.

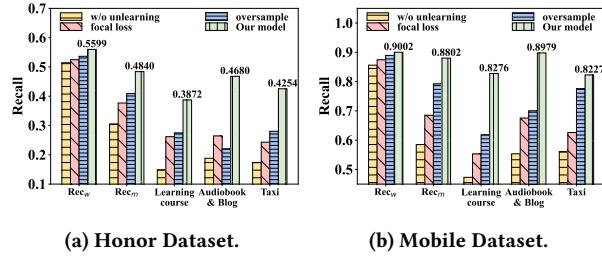
The device model results, shown in Table 3, indicate that the full-parameter tuning method achieves higher performance. However, its large parameter size makes it challenging to implement on the device side. The partial parameter method struggles to transfer individual preferences from population-level common preferences, resulting in lower performance. Although LGBM has a smaller parameter count, its stability is inferior, leading to decreased accuracy.

- Effectiveness of adaptive unlearning.

To demonstrate the effectiveness of adaptive unlearning in enhancing the accuracy of long-tail intents, we compared it with oversampling methods and focal loss [19]. We focused on the three

Table 3: Prediction performance using different model structures on the device side.

Model	Honor		Mobile		Infer speed	Params (Trainable)
	Prec _w	Prec _m	Prec _w	Prec _m		
FP	0.5496	0.4768	0.8803	0.8524	8.9ms	138M(138M)
PP	0.5217	0.4528	0.8659	0.8322	8.9ms	138M (39M)
LGBM	0.4742	0.3586	0.7943	0.6485	15us	8.8K(8.8K)
Our	0.5374	0.4693	0.8715	0.8449	2.85ms	10.86M(10.86M)

**Figure 8: effectiveness of adaptive unlearning strategy compared with other class imbalance handling methods.**

intents with the smallest proportions in the datasets and evaluated recall, which effectively reflects the model's performance in identifying long-tail intents.

The results, shown in Figure 8, indicate although focal loss and oversampling methods show some improvement, their Rec_w and Rec_m still differ, indicating they fail to address the deviation caused by the disparity in intention distribution between population and individual levels. Through adaptive unlearning, the model gradually overcomes biases towards these long-tail intents in population-level tuning, resulting in significant improvements in precision and recall.

4.6 Practicability Study

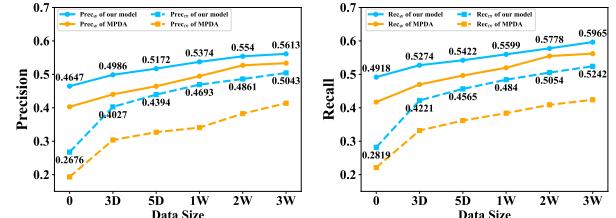
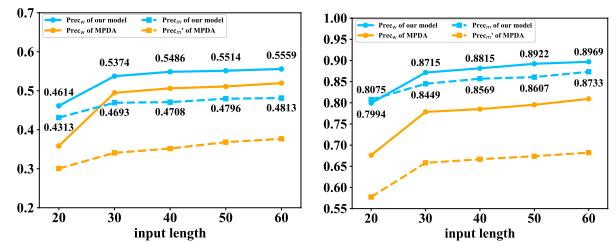
• Sensitivity of individual data scale in individual-level tuning.

In the individual-level tuning stage, particularly on the device side, there are limitations in storage and computing resources. To investigate the impact of data size, we conducted experiments by varying the data size in individual-level tuning and compared it with the second-best model (MPDA).

The results, shown in Figure 9, indicate that increasing the dataset size leads to marginal performance enhancements across all models. This trend highlights our model's capability to capture user behavior preferences. However, larger datasets significantly increase demands for computing power and storage resources. Therefore, to strike a balance between model effectiveness and computational efficiency, we selected a one-week dataset size.

• sensitivity of event sequence length.

To investigate the influence of the event sequence length, we conduct experiments by changing the input length of the historical series, and compared with the second-best model (MPDA). The results, illustrated in Figure 10, show a slight improvement in performance across all models with increasing input length. This trend highlights our model's ability to capture long-term dependencies.

**(a) Precision of Honor Dataset. (b) Recall of Honor Dataset.**
Figure 9: Influence of individual data size on performance.**(a) Precision of Honor Dataset. (b) Precision of Mobile Dataset.**
Figure 10: Influence of event seq. length on performance.

However, longer input lengths substantially increase computational demands. Thus, to balance model performance and computational efficiency, we selected an input length of 30.

5 RELATED WORKS

5.1 User Intent Prediction

User intent prediction model, a recommendation system, emphasize modeling user-event interaction sequences. Recent works [27, 41] integrate transformers into various models. Yang et al. [48] introduce intent-aware ranking with transformers, incorporating intent-aware utterance attention. Meanwhile, Wang et al. [42] propose a masked-field framework for distinct representations per intent.

Recent advancements [15, 43] focus on leveraging Graph Neural Networks (GNNs) [31] to model intent transitions and spatio-temporal features. Li et al. [18] introduce AutoIntent, featuring disentangled intent encoders and intent discovery decoders. They construct dual hyper-graphs to capture relationships and intent features. Ping et al. [26] propose an intent detection and prediction system combining human expert knowledge and consumption information to capture user preferences and context. With the rise of large language models (LLM), researchers have begun to use LLM agents to simulate behavioral intents [6, 50]. Shao et al. [32] develop an LLM workflow named Chain-of-Planned Behaviour for mobility behavior generation, which reflects the important spatial-temporal dynamics of human activities. To solve the problem of insufficient user data, Yuan et al. [51, 52] motivated Maslow's need theory, propose a knowledge-driven simulation framework based on generative adversarial imitation learning.

However, the above methods only cover some scenes in daily life resulting in the user behaviors being discontinuous and incomplete. Therefore, they can not deeply explore the user's common patterns and individual differences behind the user behavior sequences.

5.2 On-device Recommendation Model

Device-side recommender systems diverge from cloud-side recommendations by transferring model processing from the cloud to the device. This paradigm encompasses three primary approaches: (1) Device-side deployment, where models are trained in the cloud and deployed directly onto devices [43]. (2) Device-side learning, where models are trained directly on devices, often employing collaborative learning [8]. (3) Device-cloud collaboration, integrating devices with cloud-based models to enhance performance [24, 40, 49, 54]. Yan et al. [47] propose MPDA, which augments the user’s local data by retrieving similar data from the cloud’s pool. Ding et al. [5] introduce a collaborative learning framework that vertically divides the base model into two submodels: a larger one for cloud-side samples and a smaller one for device-side data, incorporating the output of the larger model.

Recent studies integrate meta-learning into recommendations to learn shared global meta-parameters to quickly adapt to individual user-specific parameters [25]. We et al. [44] propose CLOVER, a comprehensive fair meta-learning framework, which introduces a multi-task adversarial learning scheme to satisfy fairness. Kim et al. [13] propose a recommendation framework based on gradient-based meta-learning that captures the imbalanced rating distribution of each user and computes adaptive loss for user-specific learning.

However, the above methods do not consider the difference in the distribution of cloud data and device data, which is not conducive to personalized learning.

5.3 Cross-domain Fine-tuning of Pretrained LM

This year we have witnessed rapid advancements in NLP foundation models, with increasing applications of LLMs in recommendation. Two main paradigms emerge: (1) Prompt tuning, where contextual tokens guide the model’s response [14]. Geng et al. [7] propose P5 first employ LLMs in a unified text-to-text approach. (2) Instruction tuning involves detailed text instructions to enhance zero-shot model performance [3]. Bao et al. [2] propose TALLRec, align LLMs with recommendations through data tuning. Wei et al. [45] present LLMRec, enhances systems via LLM-based graph augmentation.

Moreover, the transformer, a fundamental component of LLM, tokenizes inputs into embeddings, endowing it with universal representation for cross-domain transfer. Lu et al. [23] illustrates that PLM enhances performance and computational efficiency in non-language downstream tasks. Tian et al. [38] offers a unified framework for diverse time series tasks, showing that PLM yields comparable performance across main time series analysis tasks. Jin et al. [11] introduce Time-LLM, a reprogramming framework for general time series forecasting, aligning time series with text prototypes to reconcile two modalities. Liu et al. [20] propose UniTime for multivariate time series forecasting, employing domain instructions and a language-TS transformer to achieve zero-shot transferability through modality alignment.

LLMs harness a rich dataset of human behaviors during training, encompassing prevalent patterns, common sense, and underlying rules, yet the application of LLMs in simulating human behavior and user intent prediction remains an unexplored territory.

6 CONCLUSION

Our research adapting PLMs into the human behavioral domain for on-device user intent prediction. We propose a population-to-individual tuning framework, which contains two main stages. In the population-level tuning stage, we leverage a PLM to capture the population-level common behavior patterns with the event reconstruction loss to enhance the event-to-intent transition pattern and obtain a lightweight predictor by model distillation. In the individual-level tuning framework, we utilize adaptive unlearning to correct the bias in long-tail intents due to the inconsistency between the intent distribution on population-level and individual-level. Finally, we use the individual user data to finetune and derive a personalized intent prediction model.

In future work, we aim to extend the number of intents and use disentanglement methods [28] to implement debiased learning to solve the problem of insufficient learning of long-tail intents. Besides, we aim to consider the semantics to enhance behavior understanding and prediction by urban knowledge graph [21, 22].

ACKNOWLEDGMENTS

This research has been supported in part by BNRist, National Key Research and Development Program of China under Grant 2022YFB3104702; in part by the National Natural Science Foundation of China under Grant 62272262 and Grant U23B2030; in part by the joint project of Honor Inc. & Tsinghua University.

REFERENCES

- [1] Sarfraz Ahmed, M Nazmul Huda, Sujan Rajbhandari, Chitta Saha, Mark Elshaw, and Stratis Kanarachos. 2019. Pedestrian and cyclist detection and intent estimation for autonomous vehicles: A survey. *Applied Sciences* 9, 11 (2019), 2353.
- [2] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. TALLRec: An Effective and Efficient Tuning Framework to Align Large Language Model with Recommendation (*RecSys ’23*). Association for Computing Machinery, New York, NY, USA, 1007–1014. <https://doi.org/10.1145/3604915.3608857>
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [4] Jiaao Chen and Dify Yang. 2023. Unlearn What You Want to Forget: Efficient Unlearning for LLMs. *ArXiv* abs/2310.20150 (2023). <https://api.semanticscholar.org/CorpusID:264828972>
- [5] Yucheng Ding, Chaoyue Niu, Fan Wu, Shaojie Tang, Chengfei Lyu, and Guihai Chen. 2023. DC-CCL: Device-Cloud Collaborative Controlled Learning for Large Vision Models. *arXiv preprint arXiv:2303.10361* (2023).
- [6] Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhiyun Zhou, Fengli Xu, and Yong Li. 2023. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *arXiv preprint arXiv:2312.11970* (2023).
- [7] Shijie Geng, Shuchang Liu, Zuhui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (P5). In *Proceedings of the 16th ACM Conference on Recommender Systems* (Seattle, WA, USA) (*RecSys ’22*). Association for Computing Machinery, New York, NY, USA, 299–315. <https://doi.org/10.1145/3523227.3546767>
- [8] Yeting Guo, Fang Liu, Zhiping Cai, Hui Zeng, Li Chen, Tongqing Zhou, and Nong Xiao. 2021. PREFER: Point-of-interest REcommendation with efficiency and privacy-preservation via Federated Edge leaRning. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 1, Article 13 (mar 2021), 25 pages. <https://doi.org/10.1145/3448099>
- [9] John A Hartigan and Manchek A Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)* 28, 1 (1979), 100–108.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2014. Distilling the knowledge in a neural network. In *Neural Information Processing Systems 2014 Workshop on Deep Learning and Representation Learning*.
- [11] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, et al. 2023. Time-llm:

Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728* (2023).

[12] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems* 30 (2017).

[13] Minchang Kim, Yongjin Yang, Jung Hyun Ryu, and Taesup Kim. 2023. Meta-Learning with Adaptive Weighted Loss for Imbalanced Cold-Start Recommendation (CIKM '23). Association for Computing Machinery, New York, NY, USA, 1077–1086. <https://doi.org/10.1145/3583780.3614965>

[14] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691* (2021).

[15] Jiayu Li, Peijie Sun, Zhefan Wang, Weizhi Ma, Yangkun Li, Min Zhang, Zhoutian Feng, and Daiyue Xue. 2023. Intent-aware Ranking Ensemble for Personalized Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (<conf-loc>, <city>Taipei</city>, <country>Taiwan</country>, <conf-loc>)* (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 1004–1013. <https://doi.org/10.1145/3539618.3591702>

[16] Tong Li, Yali Fan, Yong Li, Sasu Tarkoma, and Pan Hui. 2021. Understanding the long-term evolution of mobile app usage. *IEEE Transactions on Mobile Computing* 22, 2 (2021), 1213–1230.

[17] Tong Li, Tong Xia, Huandong Wang, Zhen Tu, Sasu Tarkoma, Zhu Han, and Pan Hui. 2022. Smartphone app usage analysis: datasets, methods, and applications. *IEEE Communications Surveys & Tutorials* 24, 2 (2022), 937–966.

[18] Yinfeng Li, Chen Gao, Xiaoyi Du, Huazhou Wei, Hengliang Luo, Depeng Jin, and Yong Li. 2022. Automatically Discovering User Consumption Intents in Meituan. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 3259–3269. <https://doi.org/10.1145/3534678.3539122>

[19] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42 (2017), 318–327. <https://api.semanticscholar.org/CorpusID:206771220>

[20] Xu Liu, Junfeng Hu, Yuan Li, Shizhe Diao, Yuxuan Liang, Bryan Hooi, and Roger Zimmermann. 2023. UniTime: A Language-Empowered Unified Model for Cross-Domain Time Series Forecasting. *arXiv preprint arXiv:2310.09751* (2023).

[21] Yu Liu, Jingtao Ding, Yanjie Fu, and Yong Li. 2023. Urbang: An urban knowledge graph system. *ACM Transactions on Intelligent Systems and Technology* 14, 4 (2023), 1–25.

[22] Yu Liu, Zhilun Zhou, Yong Li, and Depeng Jin. 2023. Urban knowledge graph aided mobile user profiling. *ACM Transactions on Knowledge Discovery from Data* 18, 1 (2023), 1–30.

[23] Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. 2022. Frozen Pre-trained Transformers as Universal Computation Engines. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 7 (Jun. 2022), 7628–7636. <https://doi.org/10.1609/aaai.v36i7.20729>

[24] Zheqi Lv, Wenqiao Zhang, Shengyu Zhang, Kun Kuang, Feng Wang, Yongwei Wang, Zhengyu Chen, Tao Shen, Hongxia Yang, Beng Chin Ooi, et al. 2023. DUET: A Tuning-Free Device-Cloud Collaborative Parameters Generation Framework for Efficient Device Model Generalization. In *Proceedings of the ACM Web Conference 2023*. 3077–3085.

[25] Xingyu Pan, Yushuo Chen, Changxin Tian, Zihan Lin, Jinpeng Wang, He Hu, and Wayne Xin Zhao. 2022. Multimodal Meta-Learning for Cold-Start Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Atlanta, GA, USA) (CIKM '22). Association for Computing Machinery, New York, NY, USA, 3421–3430. <https://doi.org/10.1145/3511808.3557101>

[26] Yukun Ping, Chen Gao, Taichi Liu, Xiaoyi Du, Hengliang Luo, Depeng Jin, and Yong Li. 2021. User Consumption Intention Prediction in Meituan. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (Virtual Event, Singapore) (KDD '21). Association for Computing Machinery, New York, NY, USA, 3472–3482. <https://doi.org/10.1145/3447548.3467178>

[27] Chaoyi Pu, Zhiang Wu, Hui Chen, Kai Xu, and Jie Cao. 2018. A Sequential Recommendation for Mobile Apps: What Will User Click Next App?. In *2018 IEEE International Conference on Web Services (ICWS)*. 243–248. <https://doi.org/10.1109/ICWS.2018.00038>

[28] Yuhang Quan, Jingtao Ding, Chen Gao, Nian Li, Lingling Yi, Depeng Jin, and Yong Li. 2023. Alleviating Video-length Effect for Micro-video Recommendation. *ACM Transactions on Information Systems* 42, 2 (2023), 1–24.

[29] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. <https://api.semanticscholar.org/CorpusID:160025533>

[30] Germans Savcisens, Tina Eliassi-Rad, Lars Kai Hansen, Laust Hvas Mortensen, Lau Lilleholt, Anna Rogers, Ingo Zettler, and Sune Lehmann. 2023. Using sequences of life-events to predict human lives. *Nature Computational Science* (2023), 1–14.

[31] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 20, 1 (2009), 61–80. <https://doi.org/10.1109/TNN.2008.2005605>

[32] Chenyang Shao, Fengli Xu, Bingbing Fan, Jingtao Ding, Yuan Yuan, Meng Wang, and Yong Li. 2024. Beyond Imitation: Generating Human Mobility from Context-aware Reasoning with Large Language Models. *arXiv preprint arXiv:2402.09836* (2024).

[33] Zezhi Shao, Zhao Zhang, Fei Wang, and Yongjun Xu. 2022. Pre-training Enhanced Spatial-temporal Graph Neural Network for Multivariate Time Series Forecasting. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Washington, DC, USA, August 14 – 18, 2022. ACM, 1567–1577.

[34] Neha Sharma, Chhavi Dhiman, and S Indu. 2022. Pedestrian intention prediction for autonomous vehicles: A comprehensive survey. *Neurocomputing* (2022).

[35] Tianhao Shi, Yang Zhang, Zhihui Xu, Chong Chen, Fuli Feng, Xiangnan He, and Qi Tian. 2023. Preliminary Study on Incremental Learning for Large Language Model-based Recommender Systems. *ArXiv abs/2312.15599* (2023). <https://api.semanticscholar.org/CorpusID:266550783>

[36] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.

[37] Jiaxi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (WSDM '18). Association for Computing Machinery, New York, NY, USA, 565–573. <https://doi.org/10.1145/3159652.3159656>

[38] Zhou Tian, Niu Peisong, Wang Xue, Sun Liang, and Jin Rong. 2023. One Fits All: Power General Time Series Analysis by Pretrained LM. In *NeurIPS*.

[39] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).

[40] Guanqun Wang, Jiaming Liu, Chenxuan Li, Junpeng Ma, Yuan Zhang, Xinyu Wei, Kevin Zhang, Maurice Chong, Ray Zhang, Yijiang Liu, et al. 2023. Cloud-Device Collaborative Learning for Multimodal Large Language Models. *arXiv preprint arXiv:2312.16279* (2023).

[41] Jianling Wang, Kaize Ding, Ziwei Zhu, and James Caverlee. 2021. Session-based Recommendation with Hypergraph Attention Networks. *ArXiv abs/2112.14266* (2021). <https://api.semanticscholar.org/CorpusID:232073844>

[42] Peng Wang, Jiang Xu, Chunyi Liu, Hao Feng, Zang Li, and Jieping Ye. 2020. Masked-field Pre-training for User Intent Prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, New York, NY, USA, 2789–2796. <https://doi.org/10.1145/3340531.3412726>

[43] Shoujin Wang, Liang Hu, Yan Wang, Quan Z. Sheng, Mehmet Orgun, and Longbing Cao. 2019. Modeling multi-purpose sessions for next-item recommendations via mixture-channel purpose routing networks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (Macao, China) (IJCAI '19). AAAI Press, 3771–3777.

[44] Tianxin Wei and Jingrui He. 2022. Comprehensive Fair Meta-learned Recommender System. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 1989–1999. <https://doi.org/10.1145/3534678.3539269>

[45] Wei Wei, Xubin Ren, Jiaxin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2023. Llmrec: Large language models with graph augmentation for recommendation. *arXiv preprint arXiv:2311.00423* (2023).

[46] Xin Xia, Junliang Yu, Qinyong Wang, Chaoqun Yang, Nguyen Quoc Viet Hung, and Hongzhi Yin. 2023. Efficient On-Device Session-Based Recommendation. *ACM Trans. Inf. Syst.* 41, 4, Article 102 (mar 2023), 24 pages. <https://doi.org/10.1145/3580364>

[47] Yikai Yan, Chaoyue Niu, Renjie Gu, Fan Wu, Shaojie Tang, Lifeng Hua, Chengfei Lyu, and Guihai Chen. 2022. On-Device Learning for Model Personalization with Large-Scale Cloud-Coordinated Domain Adaption. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2180–2190.

[48] Liu Yang, Minghui Qiu, Chen Qu, Cen Chen, Jiafeng Guo, Yongfeng Zhang, W. Bruce Croft, and Haiqing Chen. 2020. IART: Intent-aware Response Ranking with Transformers in Information-seeking Conversation Systems. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 2592–2598. <https://doi.org/10.1145/3366423.3380011>

[49] Jiangchao Yao, Feng Wang, Kunyang Jia, Bo Han, Jingren Zhou, and Hongxia Yang. 2021. Device-cloud collaborative learning for recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3865–3874.

[50] Yuan Yuan, Jingtao Ding, Jie Feng, Depeng Jin, and Yong Li. 2024. UniST: A Prompt-Empowered Universal Model for Urban Spatio-Temporal Prediction. *arXiv preprint arXiv:2402.11838* (2024).

[51] Yuan Yuan, Jingtao Ding, Huandong Wang, and Depeng Jin. 2024. Generating daily activities with need dynamics. *ACM Transactions on Intelligent Systems and Technology* 15, 2 (2024), 1–28.

[52] Yuan Yuan, Huandong Wang, Jingtao Ding, Depeng Jin, and Yong Li. 2023. Learning to simulate daily activities via modeling dynamic human needs. In *Proceedings of the ACM Web Conference 2023*. 906–916.

[53] Junjie Zhang, Ruobing Xie, Yupeng Hou, Wayne Xin Zhao, Leyu Lin, and Jirong Wen. 2023. Recommendation as Instruction Following: A Large Language Model Empowered Recommendation Approach. *ArXiv* abs/2305.07001 (2023). <https://api.semanticscholar.org/CorpusID:258615776>

[54] Ruiqi Zheng, Liang Qu, Tong Chen, Lizhen Cui, Yuhui Shi, and Hongzhi Yin. 2024. Decentralized Collaborative Learning with Adaptive Reference Data for On-Device POI Recommendation. *arXiv preprint arXiv:2401.13448* (2024).

A DETAILS OF MODEL DISTILLATION

In the distillation process, the teacher model produces soft targets, essentially probability distributions across intents. The student model is then trained to mimic these soft targets, rather than the actual outputs of the teacher model. This allows the student model to learn from the teacher model's knowledge without needing to replicate the same level of computational complexity [10]. Figure 11 shows the specific process of the model distilling.

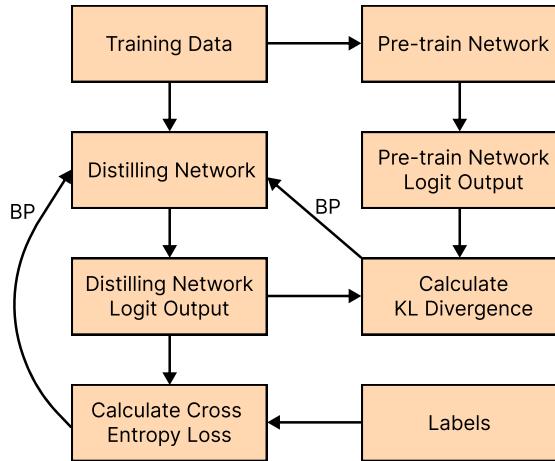


Figure 11: The process of model distilling.

B IMPLEMENTATION DETAILS FOR REPRODUCIBILITY.

Here we provide detailed values of the hyperparameters in Table 4 for reproducibility.

Table 4: Values of the hyperparameters

Hyperparameters		Value
Population-level tuning	Pretrain LM layer of transformer block	GPT2-small
	dimension of transformer block	12
	α	0.5
	ϵ	1
Individual-level tuning	layer of transformer block	4
	dimension of transformer block	768
	λ	1

C DETAILS OF BASELINES.

Here we introduce the details of each baseline.

- **CLOVER** [44]. CLOVER is a meta-learned recommendation models which consider three kinds of fairness: individual fairness, counterfactual fairness and group fairness through an adversarial learning method.
- **MetaBert4Rec** [13]. MetaBert4Rec is a sequential recommendation framework based on gradient-based meta-learning to capture the imbalanced rating distribution of each user.
- **P5** [7]. P5 is the first work to propose a unified paradigm that integrates various recommendation-related tasks into a shared conditional language generation framework.
- **InstructRec** [53]. InstructRec considers that preferences or needs can be expressed in natural language descriptions so that LLM can understand and execute the instruction for fulfilling the recommendation task.
- **LSAT** [35]. LSAT utilize two adaptation LoRA modules to learn long-term and short-term user preferences separately and then integrates them to merge the different types of preferences.
- **One fits All (OFA)** [38] OFA is a unified framework that uses a frozen pre-train language model to investigate cross-modality knowledge transfer for time series forecasting tasks.
- **TallRec** [2] TallRec is an efficient Tuning framework for Aligning LLMs with Recommendations, which structures the recommendation data as instructions and tunes the LLM via an additional instruction tuning process.
- **EODRec** [46] EODRec is an ultra-compact efficient on-device session-based recommendation that integrates discrete compositional code learning into recommendation systems to compress an item embedding table.
- **MPDA** [47] MPDA is a new device-cloud collaborative learning framework whose general idea is to retrieve some similar data from the cloud's global pool to augment the user's local data as the target domain. We choose the OFA model as the backbone.

D DETAILS OF METRICS.

we employ five widely used metrics: weighted precision ($Prec_w$), weighted recall (Rec_w), macro precision ($Prec_m$), macro recall (Rec_m), and NDCG(N). The calculation of each metric is as follows. The formula for $Prec_w$:

$$Prec_w = \frac{\sum_{c \in C} (TP_c + FP_c) \cdot Precision_c}{\sum_{c \in C} (TP_c + FP_c)} \quad (12)$$

The formula for Rec_w :

$$Rec_w = \frac{\sum_{c \in C} (TP_c + FN_c) \cdot Recall_c}{\sum_{c \in C} (TP_c + FN_c)} \quad (13)$$

The formula for $Prec_m$:

$$Prec_m = \frac{1}{|C|} \sum_{c \in C} \frac{TP_c}{TP_c + FP_c} \quad (14)$$

The formula for Rec_m :

$$Rec_m = \frac{1}{|C|} \sum_{c \in C} \frac{TP_c}{TP_c + FN_c} \quad (15)$$

Where $|C|$ represents the total number of classes, True Positives (TP_c) denotes the number of samples correctly classified as class c , False Positives (FP_c) represents the number of samples incorrectly

classified as class c , and False Negatives (FN_c) stands for the number of samples incorrectly classified as other classes instead of class c . And Precision $_c$ and Recall $_c$ respectively refer to the precision and recall of class c .

The formula for $N@k$:

$$N@k = \frac{\sum_{i=1}^K \frac{2^{rel_i} - 1}{\log_2(i+1)}}{\sum_{j=1}^{|REL_K|} \frac{rel_j - 1}{\log_2(j+1)}} \quad (16)$$

where rel_i means the graded relevance of the result at position i , and $|REL_K|$ means the list of predictions in the result ranking list up to position K .

E DETAILS OF THE TREE MODEL

We employ LightGBM [12] as our tree model, a method widely used in competitions. The hyperparameters we use for the model are shown in the table below.

Table 5: Hyperparameters Setting of LightGBM

Hyperparameters	Value
objective	multiclass
boosting	gbdt
num class	18
num iterations	2000
num leaves	32
max depth	-1
min data in leaf	20
feature fraction	1
early stopping round	75
λ_{l1}	0
λ_{l2}	0
random state	42

Please note that if the size of the dataset is less than 200, we will reduce the complexity of LightGBM by setting max depth=3, num leaves=3, $\lambda_{l1}=1$, $\lambda_{l2}=1$.

Our feature Set is shown below:

- (1) Output probabilities of all categories from GPT-2.
- (2) Position features (whether in the top 10 frequent locations).
- (3) Current hour, current day of the week, current timestamp, and indicators for morning/afternoon/evening and week-day/weekend.
- (4) For each category (illustrated by event e), time difference T_0 between the current time and the time of the last occurrence of event e , time difference T_1 between the time of the last occurrence of event e and the time of the second-to-last occurrence of event e . Given $x_i = (u_i, l_i, t_i, e_i)$, suppose event e occurred n times before the current event t_i . The time of the n_{th} occurrence of event e is denoted as t_{en} . The explanation of the mathematical formula for the time difference is as follows

$$T_0 = t_i - t_{en} \quad (17)$$

$$T_1 = t_{en} - t_{e(n-1)} \quad (18)$$