

MFNP: A Meta-optimized Model for Few-shot Next POI Recommendation

Huimin Sun¹, Jiajie Xu^{1*}, Kai Zheng^{2*}, Pengpeng Zhao^{1*},
Pingfu Chao¹ and Xiaofang Zhou³

¹School of Computer Science and Technology, Soochow University

²University of Electronic Science and Technology of China

³The Hong Kong University of Science and Technology

hmsun25@stu.suda.edu.cn, {xujj, ppzhao}@suda.edu.cn, zhengkai@uestc.edu.cn,
hellisk@outlook.com, zxf@cse.ust.hk

Abstract

Next Point-of-Interest (POI) recommendation is of great value for location-based services. Existing solutions mainly rely on extensive observed data and are brittle to users with few interactions. Unfortunately, the problem of few-shot next POI recommendation has not been well studied yet. In this paper, we propose a novel meta-optimized model MFNP, which can rapidly adapt to users with few check-in records. Towards the cold-start problem, it seamlessly integrates carefully designed user-specific and region-specific tasks in meta-learning, such that region-aware user preferences can be captured via a rational fusion of region-independent personal preferences and region-dependent crowd preferences. In modelling region-dependent crowd preferences, a cluster-based adaptive network is adopted to capture shared preferences from similar users for knowledge transfer. Experimental results on two real-world datasets show that our model outperforms the state-of-the-art methods on next POI recommendation for cold-start users.

1 Introduction

POI recommendation has become an important sub-field of recommender system [Feng *et al.*, 2015; Liu *et al.*, 2016]. As a natural extension of general POI recommendation, next POI recommendation [Cheng *et al.*, 2013] aims to predict which POI in a target region a user will go next. Compared with the general POI recommendation, next POI recommendation focuses more on exploiting sequential patterns hidden in the check-in sequence to fulfill users' current demands.

In the literature, existing methods adopt sequential analysis models like improved RNN [Feng *et al.*, 2018], LSTM [Zhao *et al.*, 2020], and GRU [Manotumruksa *et al.*, 2018] to capture dynamic user preferences for Next

POI recommendation. However, the check-in records of users may be sparse and insufficient to support effective model training, which impedes effective recommendation of above methods for cold-start users [Luo *et al.*, 2020]. To help them identify potentially interested POIs, some approaches solve the sparsity problem of cold-start user by utilizing auxiliary information, such as geographical influence [Chang *et al.*, 2018], cross-domain interactions [Zhang *et al.*, 2020] and social correlations [Cho *et al.*, 2011]. However, these approaches may recommend same POIs for users with similar context, and neglect users' real interests accordingly. It thus calls for few-shot models that can fully capture the preferences of cold-start users.

Inspired by recent advances of meta-learning in few-shot learning, some studies [Lee *et al.*, 2019; Dong *et al.*, 2020] are proposed to adopt it to solve the cold start problem in general recommendation. Meta-learning learns the experience between similar tasks for quick adaptation to new tasks with only a few training data. The recommendation for each user is usually regarded as a learning task, and a global parameter with strong generalization ability is learned to guide model parameter initialization of each task. The initialized parameters are then turned by target user's check-in data, until task-specific parameters are derived for personalized recommendation. In this way, the preferences of cold-start users can be effectively modelled through the experience of modelling preferences of users in warm state. Since meta-learning has been proven to be successful for general recommendation in cold-start scenarios [Bharadhwaj, 2019; Lu *et al.*, 2020], it is a great opportunity for few-shot next POI recommendation.

Different from general recommendation, next POI recommendation requires to consider the phenomenon of user interest drift [Yin *et al.*, 2016], since a user may have different preferences in different regions. Existing meta-learning based methods derive personal preferences through user-specific setting [Lee *et al.*, 2019; Dong *et al.*, 2020], which neglect the above phenomenon. Unfortunately, modelling personal preferences in region granularity faces severe data sparsity problem. To al-

*Corresponding author

leviate this sparsity problem, crowd preferences hidden in all check-ins in this region can be integrated. Since similar users always follow similar behaviour patterns, the crowd preferences should be obtained from similar users, instead of all users to avoid divergence. Besides, it is also challenging to design meta-learning tasks that can accurately model region-dependent user preferences by transferring the knowledge of user and crowd preferences via adaptive fusion.

In this paper, we investigate few-shot next POI recommendation while considering the phenomenon of user interest drift, and propose a Meta-optimized model for Few-shot Next POI recommendation (MFNP in short). In MFNP, we build a meta-learner for each user to provide personalized model, based on their own check-in records. Aside from user-specific tasks, our model also consists of region-specific tasks to capture region-dependent crowd preferences for the target region. To better transfer shared preferences cross similar users, we group all users by clustering first and adopt an adaptive network to balance the importance of crowd preferences of different user groups. At last, to address data sparsity, a fusion method is proposed to adaptively fuse personal preferences and region-dependent crowd preferences, such that region-dependent user preferences can be modelled via knowledge transfer. The main contributions of this work are as follows:

- We propose a novel meta-optimized model, which is the first to utilize meta-learning for few-shot next POI recommendation. By considering user interest drift, it integrates carefully designed user-specific and region-specific tasks, such that region-level user preferences can be captured by the transferred knowledge of personal and crowd preferences via adaptive fusion.
- We propose a cluster-based adaptive network to model more meaningful region-dependent crowd preferences. By merging different types of crowd patterns on clustered users, it ensures shared preferences to be shared among similar users, thus avoids the divergence among different types of users.
- Experimental results on two real datasets show that our model outperforms state-of-the-art methods on next POI recommendation for cold-start users.

2 Related Work

2.1 Next POI Recommendation

Researches on next-POI recommendation aim to exploit sequential patterns of users. Some early works [Cheng *et al.*, 2013] utilize personalized Markov chains to capture the sequential correlation between POIs. Recently, some efforts take advantage sequential deep neural models, and design customized gates in the RNN [Liu *et al.*, 2016; Feng *et al.*, 2018], LSTM [Sun *et al.*, 2020] and GRU [Manotumruksa *et al.*, 2018] to make full use of spatio-temporal information. For example, Time-LSTM [Zhu *et al.*, 2017] designs a temporal gate for

integrating temporal contexts, and an additional spatial gate is been equipped in STGN [Zhao *et al.*, 2020] to learn from spatial contexts. Since users must physically visit POIs in the real world, most models are suffered from the sparse data of users’ check-ins [Feng *et al.*, 2018]. To help cold-start users identify potentially interested POIs, some approaches utilize ancillary information. Some works try to combine the social influence [Cho *et al.*, 2011] or contextual influence [He *et al.*, 2016] with human mobility patterns.

However, these approaches may recommend same POIs for users with similar content and neglect users’ real interests. Few-shot learning models are thus needed to fully capture the preferences of cold-start users.

2.2 Meta-learning for Cold-start Recommendation

As meta-learning has proven to be successful in the few-shot learning field, recent works introduce it into the recommender system to overcome the cold-start problem. Following the settings of few-shot learning, the early works [Li *et al.*, 2019] regard each user as a task, and formulate cold-start recommendation as a few-shot learning task with user profiles. They mainly adopt optimization-based meta-learning approach such as model-agnostic meta-learning for model training. Moreover, these methods are designed for the general recommendation along with abundant auxiliary information (e.g. user profiles, item attributes). MeLU [Lee *et al.*, 2019] identifies reliable evidence candidates based on the users’ profile and the attributes of items. Also relying on the auxiliary information of users and items, MAMO [Dong *et al.*, 2020] extends two designed memory matrices to provide personalized initialization for each user.

However, the above methods neglect user interest drift and fail to model user preference in region granularity. To this end, this paper proposes a novel meta-optimized approach to effectively model region-dependent user preferences for few-shot next POI recommendation.

3 Definition

Let $U=\{u_1, u_2, \dots, u_n\}$ denote a set of users and $V=\{v_1, v_2, \dots, v_n\}$ be a set of POIs. Follow the work in [Yin *et al.*, 2016], we use K-means to cluster all the POIs into several regions according to the geographic distances, where regions can be denoted as $R=\{r_1, r_2, \dots, r_n\}$. For each POI v_i , we use l_i to denote its corresponding geographical attribute in terms of longitude and latitude coordinates, and r_i to denote its region for each POI v_i ($v_i \in V$).

Definition of check-in sequence A user u ’s check-in sequence is represented by an ordered list: $S(u) = \{v_1, v_2, \dots, v_j, v_k, \dots, v_{|S_u|}\}$, $v_{|S_u|}$ denotes the current POI. In this list, $j < k$ indicates that behavior v_j is occurred before v_k .

Definition of region’s trajectory set We define user u ’s check-ins in a region r as a trajectory: $s(u)_r =$

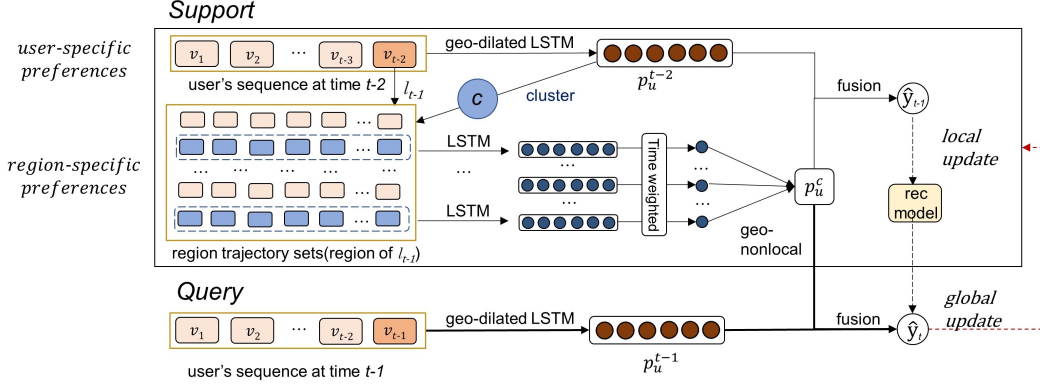


Figure 1: The Overall Architecture of MFNP Model. We split the check-in sequence of each user as support set and query set. After initializing the local parameters with the meta parameters, we model user’s preferences through the base recommender system. In the user-specific preference modelling module, we model user’s region-independent preferences according to their personal check-in records with the geo-detailed LSTM. In the region-specific preference modelling module, we first locate the corresponding user group where the user belongs to. Then we learn region-dependent crowd preferences from the region-specific check-ins and merge them by the temporal and spatial context. Finally, we locally train the recommender system with the support set, and globally update the meta parameters with the query set.

$\{v_1, v_2, \dots, v_i, \dots, v_{|s_u|}\}$, where each POI v_i is located in r . A region r ’s trajectory set is represented by a set: $s_r = \{s(u_1), s(u_2), \dots, s(u_j), \dots, s(u_{|s_r|})\}$, which is collected from all the trajectories that occurred in r . $|s_r|$ denotes the number of visitors who visited region r .

Problem Formalization Given a user $u \in U$ and a set of POIs V , with the user’s historical check-in sequence $S_u = \{v_1, v_2, v_{t-1}\}$ where v_{t-1} is the most recent POI that u has visited, next POI recommendation aims to recommend the top- k POIs that user u may be interested at a future time t .

4 Our Model

In this section, we propose a meta-optimization framework MFNP to support accurate next POI recommendation for cold-start users. Figure 1 shows the architecture of the model. We introduce the recommender model and meta-optimization learner in sections 4.1 and 4.2 respectively.

4.1 Recommender Model

The proposed base model for recommending next POI contains three modules: a user-specific module to capture personal preferences in a region-dependent way from users’ personal check-in sequences; a region-specific module to capture the similar users’ region-dependent crowd preferences in the same region; a prediction module to infer the interested POIs for users.

User-specific Preference Modelling

Inspired by previous solutions [Lee *et al.*, 2019; Li *et al.*, 2019], we first regard each user as a learning task to capture their region-independent personal preferences using their own check-in data without considering the specific region.

Understanding users’ instant demands and modelling their short-term preferences play an important role in the next POI recommendation. We represent users’ personal preferences with sequential models to explore their behaviour patterns from their complete check-in sequences. Given the check-in sequence S_u of user u , we first learn the temporal dependencies by a standard LSTM:

$$h_{t-1} = LSTM(\mathbf{x}_{t-1}, h_{t-2}), (1 \leq t \leq |S_u|) \quad (1)$$

where h_{t-1} and h_{t-2} is the hidden state of LSTM at $t-1$ and $t-2$ respectively, \mathbf{x}_{t-1} is the embedding vector for POI $v_{t-1} \in S_u$.

Since RNN-based methods strictly follow the temporal order, they can not capture the spatial dependencies from the visited POIs which are geographically scattered in the real world. Therefore, we apply the geo-dilated LSTM scheme [Sun *et al.*, 2020] to explore the geographical relations among non-consecutive POIs. The geo-dilated LSTM picks POIs from the check-in sequences as the input with different skip lengths, which determined by the geographical relevance. For example, given a check-in sequence $S_n = \{v_1, v_2, v_3, v_4, v_5\}$, v_3 has two preceding POIs in the temporal sequence: v_1 and v_2 . Since the geographical distance between v_1 and v_3 is shorter than the distance between v_1 and v_2 , there exists a geo-dilated sequences $\{v_1, v_3\}$. After finding the path from v_1 to v_5 , we can construct the input set $S_n^{geo} = \{\{v_1, v_3\}, \{v_3, v_5\}\}$. Then the preferences of users which conclude the spatial dependencies can be learned by the dilated LSTM scheme with the input S_n^{geo} :

$$h'_{t-1} = LSTM(\mathbf{x}_{t-1}, h'_{t-\delta}), (1 \leq t \leq |S_u|) \quad (2)$$

where δ is the skip length which will be automatically determined on geographical factors, $h'_{t-\delta}$ is computed from the last sequence $\{x_\delta, x_{t-1}\} \in S_n^{geo}$.

At time $t - 1$, the final representation of user’s region-independent preference is an average of the latent representations h_{t-1} and h'_{t-1} :

$$p_u^{t-1} = \frac{h_{t-1} + h'_{t-1}}{2} \quad (3)$$

Region-specific Preference Modelling

Different from the general recommendation, users’ preferences may be region-varying in the POI recommendation. However, due to the sparse check-in data, it is impossible for us to model users’ region-dependent preferences in each region separately. As the rapid development of social media and the Internet, people are more easily influenced by others’ successful experiences, resulting in that similar users may follow the similar behaviour pattern [Yin *et al.*, 2016]. This allows us to model region-dependent crowd preferences from the similar people in the same region.

To achieve this, we propose a cluster-based adaptive network to locate the user group where the user belongs to, so that region-dependent crowd preferences can be modelled more accurately. Specifically, we cluster users into different user groups based on the representation P_u^{t-1} obtained by Eq.(3) by HDBSCAN algorithm [McInnes *et al.*, 2017]:

$$c_u = HDBSCAN(p_u^{t-1}) \quad (4)$$

where c_u is the user group of u .

Given the region r where the recent POI v_{t-1} is located in, we can filter the trajectories which are generated by users which belong to c_u . All trajectories are sorted in temporal order $s_r = \{s_1, s_2, \dots, s_m, \dots, s_n\}$. The information of all POIs in each trajectory s_m is encoded by a LSTM layer:

$$h_t = LSTM(\mathbf{x}_t, h_{t-1}), (1 \leq t \leq |s_m|) \quad (5)$$

where h_{t-1} is the hidden state of LSTM, \mathbf{x}_t is the embedding vector for t -th POI $v_t \in s_m$.

Instead of a naive combination of all the corresponding trajectories, we have to incorporate the temporal factors with users’ personal trajectories. This is due to the fact that the popularity of POIs among the crowd is changing over time. For example, people may choose fast food on weekdays at lunchtime and seek for cinema on weekends night. Therefore, we apply a time weighted operation to incorporate such information to capture the time-sensitive property. We map one week into 48 time slots (24 slots for hours on weekdays and 24 slots for hours on weekends). For each slot o , we construct a POI set $G_o = \{v_1, v_2, \dots, v_{|G_o|}\}$, where $v \in V$ is a POI visited by at least one user in time slot o . We calculate the temporal similarity $\tau_{i,j}$ between the i -th and j -th time slots by Jaccard index:

$$\tau_{i,j} = \frac{|G_i \cap G_j|}{|G_i \cup G_j|} \quad (6)$$

Then, we can derive a sequence of trajectory time slots according to the check-in time of all the POIs in

s_m . The time slots sequences can be represented as $\{p_1, p_2, \dots, p_{|s_m|}\}$, where $p \in \{1, 2, \dots, 48\}$. Given the target user’s current time slot o , the sequence-level representation \mathbf{s}_m for s_m is generated as follows:

$$\mathbf{s}_m = \sum_{t=1}^{|s_m|} w_t \mathbf{h}_t, w_t = \frac{\exp(\tau_{o,p_t})}{\sum_{j=1}^{|s_m|} \exp(\tau_{o,p_j})} \quad (7)$$

where τ_{o,p_t} is the temporal similarity between the current time slot o and the time slot of the j -th visited POI in s_m . Besides, we use the similar way to generate the sequence-level representation \mathbf{S}_u for S_u , while the time-weighted operation is replaced by an average pooling.

After learning the representations $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n\}$ for all trajectories in $s_r = \{s_1, s_2, \dots, s_n\}$, we can derive region-dependent crowd preference \mathbf{p}_u^r with respect to the user’s personal trajectory S_u using a geo-nonlocal operation to integrate spatial factors:

$$\mathbf{p}_u^r = \frac{1}{C(s)} \sum_m^{|s_r|} f(\mathbf{S}_u, \mathbf{s}_m) g(\mathbf{s}_m) \quad (8)$$

where $C(S) = \sum_m^{|s_r|} \exp\left(\frac{1}{d_{u,m}} \mathbf{S}_u^\top \mathbf{s}_m\right)$ is the normalization factor by taking the user’s most recently visited location into account. $d_{u,m}$ is the geographical distance between the location l_{t-1} of the current POI v_{t-1} and the location l_{s_m} of trajectory s_m , where l_{s_m} is defined via the average pooling operation for all the visited POIs in the trajectory s_m . The pairwise function $f(\cdot)$ calculates an affinity score between the user’s personal trajectory S_u and the region-based trajectories s_m , and $g(\cdot)$ generates the representation for \mathbf{s}_m . Specifically, in our solution, $f(\mathbf{S}_u, \mathbf{s}_m)$ and $g(\mathbf{s}_m)$ are defined as:

$$f(\mathbf{S}_u, \mathbf{s}_m) = \exp\left(\frac{1}{d_{u,m}} \mathbf{S}_u^\top \mathbf{s}_m\right) \quad (9)$$

$$g(\mathbf{s}_m) = \mathbf{W}_r \mathbf{s}_m \quad (10)$$

where \mathbf{S}_u and \mathbf{s}_m are the representations of user’s personal check-in sequence S_u and region-based trajectory s_m , respectively, while \mathbf{W}_r is a trainable projection weight matrix.

Prediction

After obtaining the representations of both personal preferences \mathbf{p}_u^{t-1} and region-dependent crowd preferences \mathbf{p}_u^r , we compute the probability distribution $\hat{\mathbf{y}}$ over all POIs in $|V|$ as:

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{W}_p (\mathbf{p}_u^{t-1} \oplus \mathbf{p}_u^r)) \quad (11)$$

where \oplus is the concatenation operation of personal preferences and region-dependent crowd preferences, \mathbf{W}_p is a trainable projection matrix for all POIs. Consequently, the POI most likely to be visited by the target user at the next time step t is the one with the largest probability.

4.2 Meta-optimized framework

Local Update

Inspired by recent advances of meta-training [Li *et al.*, 2019; Dong *et al.*, 2020], we initialize the recommender

parameters from the global initial values. At the beginning of the training process, we randomly initialize the global parameters. For each user, we have a support set and a query set. During the local update phase (i.e. learning on the support set), we first initialize the local recommender parameters $\theta_u, \theta_g, \theta_r$ as the meta parameters ϕ_u, ϕ_g, ϕ_r , where $\theta_u, \theta_g, \theta_r$ represent the parameters of the user-specific preference module, region-specific preferences module and the prediction module, respectively. Calculated by the Eq.(1-3), we can obtain region-independent personal preferences. After locating the corresponding user group, we calculate the region-dependent crowd preferences by Eq.(5-10). The prediction is calculated through the above two preferences by Eq.(11). The optimization goal of local training is to minimize the loss of the recommendation for a single user, the parameters can be locally updated as follows:

$$\theta_* \leftarrow \theta_* - \rho \cdot \nabla_{\theta_*} \mathcal{L}(y, \hat{y}) \quad (12)$$

Global Update

The aim of the meta optimization is to minimize the expected loss on the query set. After the local training on the support set, we update the global parameters according to the loss on query sets. Suppose the recommender model is denoted as $\mathcal{R}_{\hat{\theta}_*}$, where $\theta = \{\theta_u, \theta_g, \theta_r\}$. After the local training on support set, we get the model $\mathcal{R}_{\hat{\theta}_*}$ with updated parameter $\hat{\theta}$. Our goal is minimizing the training loss for users $u \in U_{train}$ on query sets for $q \in Q$. Then, the global parameters are updated by:

$$\phi_* \leftarrow \phi_* - \lambda \sum_{u \in U_{train}} \sum_{q \in Q} \nabla \mathcal{L} \left(\mathcal{R}_{\hat{\theta}_*} \right) \quad (13)$$

5 Experiment

5.1 Dataset

We conduct experiments on two public LBSNs datasets, namely Foursquare [Yang *et al.*, 2015] and Gowalla [Yin *et al.*, 2013]. The statistics of both datasets are shown in Table 1. The Foursquare dataset contains the POI check-in records created by users living in New York from February 2009 to October 2010, while Gowalla contains world-wide check-ins from February 2009 to October 2010. For both datasets, we eliminate unpopular POIs that are visited by less than 3 users and trim the number of check-in records for each user as with 10 records to force the model to learn from few samples. Each user is regarded as a sample in the dataset. We randomly separate the users into training and testing users with ratio 80:20. Especially, for all the meta-optimized methods, the check-in sequence of a single user are further divided into support set and query set by employing the data augmentation strategy.

5.2 Compared methods and Settings

We adopt the following baselines for recommender systems, where MeLU and MAMO are based on meta-optimization for general recommendation, and the other methods are sequential-based model for next POI recommendation.

Table 1: Dataset description

Datasets	Users	POIs	Check-ins
Foursquare	824	38,336	227,428
Gowalla	18,737	32,510	1,278,274

- **MeLU** [Lee *et al.*, 2019] obtains auxiliary information from the user profile and item attributes and optimized the recommender model which composed of a fully connected layer as meta-learning.
- **MAMO** [Dong *et al.*, 2020] extends the MeLU with two designed memories to provide personalized initialization for each user.
- **STRNN** [Liu *et al.*, 2016] applies time-specific transition matrices and distance-specific transition matrices in RNN to integrate spatial and temporal contexts.
- **T-LSTM** [Zhu *et al.*, 2017] equips time gates with LSTM for modelling the temporal interval between check-ins from users’ sequential actions.
- **STGCN** [Zhao *et al.*, 2020] extends a spatial gate on the T-LSTM network architecture to model spatial intervals between check-ins for the next POI recommendation.
- **LSTPM** [Sun *et al.*, 2020] develops a context-aware non-local network to model long-term preferences and a geo-dilated LSTM to model short-term preferences, which considers both spatial and temporal factors.

Following the work in [Feng *et al.*, 2018], we set the dimension of embeddings and the hidden states to 500 for all deep learning-based methods. The K of region clusters is set to 30, while the number of cluster of users is set to 6. All the parameters in our model are optimized using the gradient descent optimization algorithm Adam with the batch size of 1 and the learning rate of 0.0001.

5.3 Evaluation

To evaluate the performance of each method for next POI recommendation, we adopt two evaluation metrics Accuracy ($Acc@K$) and Normalized Discounted Cumulative Gain (NDCG@K). $Acc@K$ measures the presence of the correct POI among the ground-truth, and NDCG@K measures the quality top- K ranking list. In this paper, we choose the popular $K = \{1, 5, 10\}$ for evaluation.

5.4 Performance

Firstly, although MeLU and MAMO can learn from few training samples, they still perform worse than other sequential models, due to the lack of sufficient auxiliary information in the next POI recommendation. Secondly, both LSTPM and STGCN are superior to other methods, since they both adaptively model the preferences of users according to the current states of users. At last, our

Table 2: Performance comparison on two datasets w.r.t. $Acc@K$ and $NDCG@K$

	Foursquare					Gowalla				
	$Acc@1$	$Acc@5$	$Acc@10$	$NDCG@5$	$NDCG@10$	$Acc@1$	$Acc@5$	$Acc@10$	$NDCG@5$	$NDCG@10$
MeLU	0.0022	0.0074	0.0108	0.0039	0.0063	0.0003	0.0019	0.0026	0.0004	0.0015
MAMO	0.0044	0.0145	0.0212	0.0058	0.0135	0.0019	0.0034	0.0052	0.0013	0.0023
STRNN	0.0087	0.0177	0.0230	0.0072	0.0140	0.0043	0.0053	0.0077	0.0021	0.0042
T-LSTM	0.0075	0.0178	0.0215	0.0079	0.0137	0.0043	0.0061	0.0068	0.0026	0.0039
LSTPM	0.0114	0.0214	0.0295	0.0129	0.0148	0.0087	0.0206	0.0238	0.0127	0.0158
STGCN	0.0158	0.0256	0.0297	0.0135	0.0145	0.0076	0.0212	0.0277	0.0145	0.0171
MFNP	0.0255	0.0379	0.0503	0.0173	0.0292	0.0152	0.0308	0.0381	0.0234	0.0258

proposed model significantly outperforms all the baseline methods on both datasets in terms of the $Acc@K$ and $NDCG@K$ metric. It indicates the effectiveness of the MFNP model, due to the way of building meta-learners on both user-specific tasks and region-specific tasks. It demonstrates that our model can explore the preferences from the other users and adapt it to the cold-start users. Moreover, the performance of Foursquare performs better than Gowalla. That is because the users' check-in in Gowalla is sparser than these in Foursquare. The performance of our model further verifies that our work can do well under the scenario of data sparsity in the next POI recommendation.

Effectiveness of Region-specific Preference

To verify whether the region-specific preference modelling module and the corresponding operation is necessary for next POI recommendation, we conduct the experiment on two variants of our MFNP model.

- **MFNP-User** This version removes the whole region-specific preference module and only capture users' region-independent preferences from the check-in sequence of themselves.
- **MFNP-Crowd** This version applies the average operation instead of the adaptive method for considering the spatial and temporal factors when combining region-dependent crowd preferences.

As shown in Table 3, we can observe that MFNP-Crowd always perform better than MFNP-User, showing that region-dependent crowd preferences are effective for the next POI recommendation for cold-start users. Since MFNP is always significantly better than MFNP-Crowd, we can draw the conclusion that considering region-dependent crowd preferences of similar users with spatial and temporal factors are helpful for providing personalized region-dependent crowd preferences for users.

Impact of the Number of User Clusters

To investigate whether the performance of our model is sensitive to the number of user clusters and explore the optimal parameters, we test the performance of MRNP by varying the number of user types. As shown in Figure 2, we can observe that the accuracy of the next POI recommendation first increases and then decrease and then becomes stable as the number of user clusters gradually

Table 3: Comparison of different variants of MFNP

	Foursquare		Gowalla	
	$Acc@1$	$Acc@5$	$Acc@1$	$Acc@5$
MFNP-User	0.0174	0.0254	0.0079	0.0175
MFNP-Crowd	0.0219	0.0294	0.0137	0.0278
MFNP-MFNP	0.0255	0.0379	0.0152	0.0308




Figure 2: Performance comparison of different numbers of user clusters.

increases. The optimal parameter setting of the number of clusters is $C=6$ for Foursquare and Gowalla.

6 Conclusion

In this paper, we investigate few-shot next POI recommendation while considering the phenomenon of user interest drift. We build a meta-learner for each user to provide personalized model according to their own check-in records. In addition, we propose a region-specific preference module to capture region-dependent crowd preferences of similar users in the same region with the target user. An adaptive method is proposed to combine these preferences with spatial and temporal factors. The experiments demonstrate that our model outperformed the state-of-the-art methods on real-world datasets.

Acknowledgement

This work is supported by the NSFC project under grant numbers (61872258, 61876117, 61772356, 61972069, 61836007, 61832017, 62072125), the major project of natural science research in Universities of Jiangsu Province under grant number 20KJA520005.

References

- [Bharadhwaj, 2019] Homanga Bharadhwaj. Meta-learning for user cold-start recommendation. In *IJCNN*, pages 1–8. IEEE, 2019.
- [Chang *et al.*, 2018] Buru Chang, Yonggyu Park, Donghyeon Park, Seongsoon Kim, and Jaewoo Kang. Content-aware hierarchical point-of-interest embedding model for successive poi recommendation. In *IJCAI*, pages 3301–3307, 2018.
- [Cheng *et al.*, 2013] Chen Cheng, Haiqin Yang, Michael R Lyu, and Irwin King. Where you like to go next: Successive point-of-interest recommendation. In *IJCAI*, 2013.
- [Cho *et al.*, 2011] Eunjoon Cho, Seth A Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *KDD*, pages 1082–1090, 2011.
- [Dong *et al.*, 2020] Manqing Dong, Feng Yuan, Lina Yao, Xiwei Xu, and Liming Zhu. Mamo: Memory-augmented meta-optimization for cold-start recommendation. In *KDD*, pages 688–697, 2020.
- [Feng *et al.*, 2015] Shanshan Feng, Xutao Li, Yifeng Zeng, Gao Cong, and Yeow Meng Chee. Personalized ranking metric embedding for next new poi recommendation. In *IJCAI*, pages 2069–2075. ACM, 2015.
- [Feng *et al.*, 2018] Jie Feng, Yong Li, Chao Zhang, Funing Sun, Fanchao Meng, Ang Guo, and Depeng Jin. Deepmove: Predicting human mobility with attentional recurrent networks. In *WWW*, pages 1459–1468, 2018.
- [He *et al.*, 2016] Jing He, Xin Li, Lejian Liao, Dandan Song, and William Cheung. Inferring a personalized next point-of-interest recommendation model with latent behavior patterns. In *AAAI*, volume 30, 2016.
- [Lee *et al.*, 2019] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. Melu: Meta-learned user preference estimator for cold-start recommendation. In *KDD*, pages 1073–1082, 2019.
- [Li *et al.*, 2019] Jingjing Li, Mengmeng Jing, Ke Lu, Lei Zhu, Yang Yang, and Zi Huang. From zero-shot learning to cold-start recommendation. In *AAAI*, volume 33, pages 4189–4196, 2019.
- [Liu *et al.*, 2016] Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. Predicting the next location: A recurrent model with spatial and temporal contexts. In *AAAI*, 2016.
- [Lu *et al.*, 2020] Yuanfu Lu, Yuan Fang, and Chuan Shi. Meta-learning on heterogeneous information networks for cold-start recommendation. In *KDD*, pages 1563–1573, 2020.
- [Luo *et al.*, 2020] Mi Luo, Fei Chen, Pengxiang Cheng, Zhenhua Dong, Xiuqiang He, Jiashi Feng, and Zhen-guo Li. Metaselector: Meta-learning for recommendation with user-level adaptive model selection. In *WWW*, pages 2507–2513, 2020.
- [Manotumruksa *et al.*, 2018] Jarana Manotumruksa, Craig Macdonald, and Iadh Ounis. A contextual attention recurrent architecture for context-aware venue recommendation. In *SIGIR*, pages 555–564, 2018.
- [McInnes *et al.*, 2017] Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11):205, 2017.
- [Sun *et al.*, 2020] Ke Sun, Tiejun Qian, Tong Chen, Yile Liang, Quoc Viet Hung Nguyen, and Hongzhi Yin. Where to go next: Modeling long-and short-term user preferences for point-of-interest recommendation. In *AAAI*, volume 34, pages 214–221, 2020.
- [Yang *et al.*, 2015] Dingqi Yang, Daqing Zhang, Vincent W. Zheng, and Zhiyong Yu. Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns. *SMC*, 45(1):129–142, 2015.
- [Yin *et al.*, 2013] Hongzhi Yin, Yizhou Sun, Bin Cui, Zhiting Hu, and Ling Chen. Lcars: a location-content-aware recommender system. In *KDD*, pages 221–229, 2013.
- [Yin *et al.*, 2016] Hongzhi Yin, Xiaofang Zhou, Bin Cui, Hao Wang, Kai Zheng, and Quoc Viet Hung Nguyen. Adapting to user interest drift for poi recommendation. *TKDE*, 28(10):2566–2581, 2016.
- [Zhang *et al.*, 2020] Hao Zhang, Siyi Wei, Xiaojiao Hu, Ying Li, and Jiajie Xu. On accurate POI recommendation via transfer learning. *DPD*, 38(3):585–599, 2020.
- [Zhao *et al.*, 2020] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Fuzhen Zhuang, Jiajie Xu, Zhixu Li, Victor S Sheng, and Xiaofang Zhou. Where to go next: A spatio-temporal gated network for next poi recommendation. *TKDE*, 2020.
- [Zhu *et al.*, 2017] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. What to do next: Modeling user behaviors by time-lstm. In *IJCAI*, volume 17, pages 3602–3608, 2017.