When Online Meets Offline: Exploring Periodicity for Travel Destination Prediction

Wanjie Tao
Alibaba Group
Hangzhou, Zhejiang, China 311121
wanjie.twj@alibaba-inc.com

Liangyue Li
Alibaba Group
Hangzhou, Zhejiang, China 311121
jackiey99@gmail.com

Chen Chen
University of Virginia
United States
chenannie45@gmail.com

Zulong Chen
Alibaba Group
Hangzhou, Zhejiang, China 311121
zulong.czl@alibaba-inc.com

Hong Wen
Alibaba Group
Hangzhou, Zhejiang, China 311121
qinggan.wh@alibaba-inc.com

ABSTRACT

Online travel platforms (OTPs), e.g., booking.com and Ctrip.com, deliver travel experiences to online users by providing travel-related products. One key problem facing OTPs is to predict users' future travel destination, which has many important applications, e.g., proactively recommending users flight tickets or hotels in the destination city. Although much progress has been made for the next POI recommendation, they are largely sub-optimal for travel destination prediction on OTPs, due to the unique characteristics exhibited from users' travel behaviors such as offline spatial-temporal periodicity and online multi-interest exploration. In this paper, we propose an online-offline periodicity-aware information gain network, OOPIN, for travel destination prediction on OTPs. The key components of the model are (1) an offline mobility pattern extractor, which extracts spatial-temporal periodicity along with the sequential dependencies from the visited city sequence; and (2) an online multi-interests exploration module that discovers destinations that the user might be interested in but not yet visited from their online interaction data. Comprehensive experiments on real-world OTP demonstrate the superior performance of the proposed model for travel destination prediction compared with state-of-the-art methods.

CCS CONCEPTS

• **Information systems → Recommender systems**.

KEYWORDS

Travel destination prediction, Neural network

ACM Reference Format:
Wanjie Tao, Liangyue Li, Chen Chen, Zulong Chen, and Hong Wen. 2022. When Online Meets Offline: Exploring Periodicity for Travel Destination Prediction. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22), July 11–15, 2022, Madrid, Spain.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3477495.3531859

SIGIR '22, July 11–15, 2022, Madrid, Spain

ACM ISBN 978-1-4503-8732-3/22/07…$15.00
https://doi.org/10.1145/3477495.3531859

1 INTRODUCTION

Online travel platforms (OTPs), e.g., bookings.com and Ctrip.com, deliver travel experiences to online users by providing travel-related products (e.g., flight tickets, hotels, and package tours). Different from regular e-commerce platforms, online travel platforms pay more attentions to temporal patterns (e.g., users usually travel on weekends and holidays) and spatial patterns (e.g., users travel to nearby cities or distant cities) exhibited by users' travel behaviors on the platform. More importantly, users' travel behaviors are quite sparse and divergent [12, 13], since travel is a low-frequency demand compared with shopping [5, 8, 16].

One key problem facing OTPs is to predict users' future travel destinations, which leads to many important applications for OTPs. For instance, the online travel app could push a notification to recommend flight tickets or hotels for the destination city. Furthermore, the predicted cities can in turn be used to provide fine-grained recommendations on featured local tours in the destination.

The closest line of research is on the next POI recommendations, which have drawn intensive studies owing to the rapid growth of location-based services. Earlier next POI recommendation algorithms mainly focus on modeling the transition probability between visited places by Markov chains [1, 4]. Moving forward, recurrent neural networks (RNNs) are extensively used to improve the recommendation quality [2, 6, 10, 17]. Recent state-of-the-art algorithms are largely based on adjusted neural network architectures with enriched spatial-temporal information [9, 11].

Although much progress has been made, the state-of-the-arts for next POI recommendation are largely sub-optimal for the travel destination prediction on OTPs. The sub-optimality is mainly due to the unique characteristics exhibited by users' travel-related behavior. First, *offline spatial-temporal periodicity:* one unique trait of users' traveling behavior is that they usually demonstrate periodic patterns in both the temporal and spatial dimensions. On the temporal dimension, users often go traveling on weekends or holidays; on the spatial dimension, users tend to visit nearby cities frequently and distant cities occasionally. Such spatial and temporal periodicity are also correlated with each other as users tend to visit surrounding places on weekends and travel afar on long holidays. Second, *online multi-interest exploration:* a user may explore the travel products (e.g., flight tickets and hotels) of several candidate cities on OTPs through search or platform recommendations before deciding the final destination. Extracting the information gain from such online exploration behavior on cities that have not been visited by the user can be beneficial for the next destination prediction.

With these in mind, we propose an online-offline periodicity-aware information gain network (OOPIN in short), for travel destination prediction. The model consists of two key components, namely, an offline mobility pattern extractor and an online multi-interests exploration module. We design a periodicity-aware GRU layer in the offline component, which aims to extract the sequential dependencies from offline mobility sequence (i.e. the sequence of the cities the user visited) with the awareness of spatial-temporal periodicity. To leverage users' online interaction behavior, in the online component, we design a distance-aware self-attention net to extract multiple user interests, and an information gain net along with the target-aware attention network to represent the destinations that the user is interested in but has not yet visited. The main contributions are summarized as follows:

- To the best of our knowledge, this is the first work to marry the offline mobility data and online behavioral data for the purpose of travel destination prediction on OTPs.
- We propose a novel online-offline periodicity-aware information gain network for the purpose of the destination prediction.
- Comprehensive experiments conducted both offline and online demonstrate the effectiveness of our model, which has been deployed online to serve real traffic of an OTP.

2 RELATED WORKS

Sequential Recommendation. Sequential recommendation is extensively used to predict users' next behavior from previously observed behavior sequences. Earlier work proposes to utilize the Markov model to approximate the transition probability between two consecutive actions [1, 4]. By the rapid growth of deep learning, GRU4Rec, an RNN-base model, is used for user action sequence recommendation [6]. SASRec [7] utilizes the attention mechanism to adaptively assign weights to previous items for next-item prediction. DSIN [3] models users' session interests and its evolving process with self-attention mechanism and Bi-LSTM. SINS [15] proposes to infer a sparse set of concepts for each user from the large concept pool and output multiple interest embeddings accordingly, which are effective for user intent and next-item prediction.

Next POI Recommendation. Next POI recommendation can be viewed as a special type of sequential recommendation problem, which considers spatial information in the sequence. In [10], an RNN-based POI recommendation model STRNN is introduced, which effectively incorporates spatial-temporal transition matrices. LSTPM [14] captures long-term preferences with a nonlocal network and short-term preferences with a geo-dilated RNN. DeepMove [2] combines the attention model with the RNN model to enable long-term periodicity learning in mobility prediction. STGN [17] models both short and long-term POI visit preferences based on LSTM with new time and distance gates, and cell state. GeoSAN [9] exploits hierarchical gridding of GPS locations for spatial discretization and uses self-attention layers for matching, without using explicit spatio-temporal interval. STAN [11] leverages the spatial-temporal-preference neighborhoods graphs' structures and explores new higher-order POI neighbors.

In general, the previous methods would suffer from sparsity issues when limited mobility data was observed. Our model can alleviate the problem by the complementary effect between offline mobility data and online behavior data.

3 PROBLEM STATEMENT

Let $\mathcal{U}$ be the set of users, and C be the set of cities. Given a user $u \in \mathcal{U}$, we can observe the sequence of the user's visited cities ordered by time as $\mathcal{V}_u = \{c_1^{off}, \ldots, c_T^{off}\}$. Besides, the user often exhibits online interaction behaviors with the travel products on the OTP app, e.g., clicking on a flight ticket or browsing a hotel. We denote the sequence of destination cities associated with the interacted travel products as $O_u = \{c_1^{on}, \ldots, c_n^{on}\}$. The problem of travel destination prediction is to predict the destination city that the user u is likely to visit for her next travel.

4 PROPOSED MODEL

4.1 Overall Framework

The overall framework of our model is illustrated in Fig. 1, which consists of two key components, namely, the offline mobility pattern extractor and the online multi-interests exploration module. Users' offline mobility data usually exhibits the spatial-temporal periodic evolution patterns, e.g., some users might go on business trips on weekdays, tour around on weekends and travel long-distance during holidays. Beyond periodicity, the sequential dependencies also exist in the offline mobility data, e.g., if a user visited a surrounding city last weekend for leisure travel, she might not visit the same city again for the coming weekend. The offline mobility pattern extractor takes the offline visited destinations sequence as input and then utilizes a periodicity-aware Gated Recurrent Unit (PGRU) layer to model the spatial-temporal periodic patterns. On the other hand, a user may plan her itinerary by browsing various travel products (e.g. flight tickets, scenic spots tickets, hotels, etc.) of different destination cities on OTPs before actual trips. Such online interaction behaviors are strong indications of the user's next travel destination and can also relieve the sparsity issue of offline mobility data. To this end, we design an online multi-interests exploration module that is able to extract multiple candidate destinations that the user might be interested in from her online behaviors prior to the actual trip. Among the candidate destinations, the ones that have not been visited by the user would be out-weighted through a target-aware information gain network.

4.2 Offline Mobility Pattern Extractor

The offline mobility pattern extractor takes the sequence of cities visited by the user $\mathcal{V}_u$ as input, which will be processed by a periodic pattern extractor to mine the spatial-temporal patterns and a periodicity-aware GRU layer to model the periodicity aware sequential evolution.

4.2.1 Periodic Pattern Extractor. The offline mobility sequence $\mathcal{V}_u$ is organized into a set of spatial-temporal mobility matrices $\mathcal{M} = \{\mathbf{M}_1, \ldots, \mathbf{M}_k\}$, where $\mathbf{M}_i$ is designed to represent mobility frequency at different temporal granularity and spatial distance. For example, in $\mathbf{M}_1$, the first dimension could represent the seven

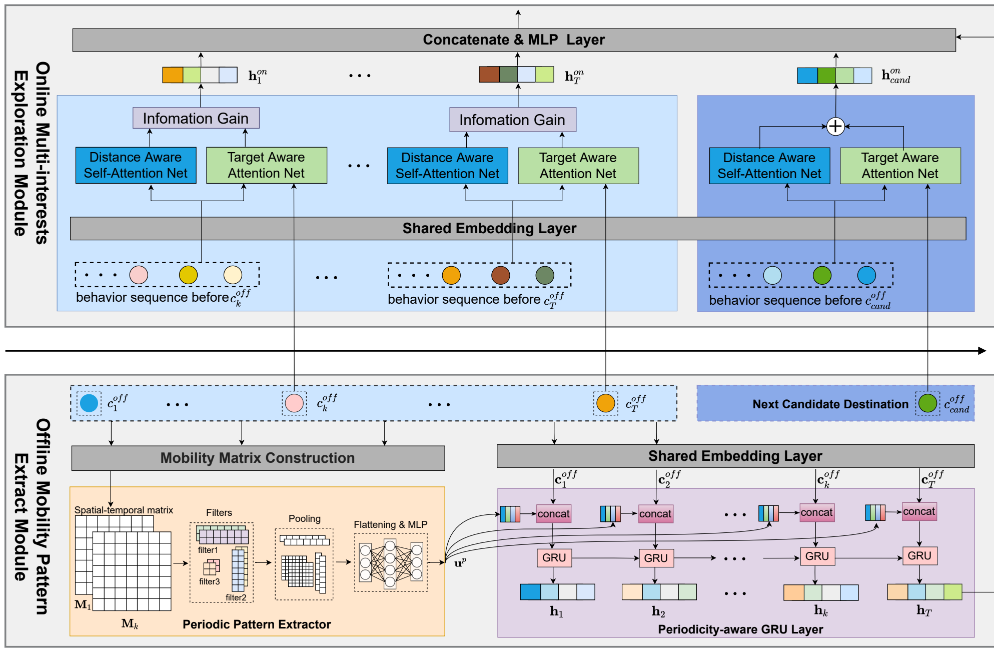

Figure 1: The overall architecture of our proposed model OOPIN.

days in a week from Monday to Sunday, and the second dimension represents the distance between the departure city and the visited city. Note that the distance is discretized into different buckets. Each cell in $\mathbf{M}_1$ represents the frequency of the user visiting a city of a particular distance on a particular day of a week. Similarly, we could design other mobility matrices with different temporal semantics, e.g., different days in a month, holidays, or other events. After this, we carefully design a set of convolutional kernels of sizes $2 \times d$, $t \times 2$, and 2×2 to extract user periodic mobility patterns in the temporal-spatial dimensions, where d is the number of distance buckets and t is the temporal dimensionality. The feature matrices after convolving the mobility matrices with the set of kernels are pooled, flattened, and concatenated, followed by an MLP layer to get the periodic pattern representation $\mathbf{u}^p$.

4.2.2 Periodicity-aware GRU Layer. Beyond the spatial-temporal periodicity, the sequential evolution also exists in the visited city sequence, where the user's recent travel behaviors might impact her next travel plan. For instance, if a user has recently returned to her hometown, it is unlikely she would return to the hometown again for her next travel. To capture such sequential dependencies, we design a periodicity-aware GRU layer, where the inputs are the periodic pattern representation $\mathbf{u}^p$ and the embeddings of visited cities in $\mathcal{V}_u$, $\{\mathbf{c}_1^{off}, \dots, \mathbf{c}_T^{off}\}$. We concatenate the periodicity signal $\mathbf{u}^p$ with $\mathbf{c}_i^{off}$ at each time step and feed the sequence into a GRU layer and get the hidden representation at the last time step $\mathbf{h}_T$.

4.3 Online Multi-interests Exploration Module

The user offline mobility data is often very sparse since traveling is a low-frequency demand. However, the user usually explores the travel products (e.g., flight tickets, hotels) of several candidate cities on OTPs. Such online interaction behavior could complement the offline mobility data for the travel destination prediction.

We split the user's online interacted city sequence O_u into multiple subsequences $\mathcal{J}_i \subseteq O_u$, where $\mathcal{J}_i$ is the subsequence of online interacted cities before and during an actual visit of a city c_i^{off}, i.e., $\mathcal{J}_i = \{c_{i1}^{on}, \dots, c_{it}^{on}\}$.

Distance-aware Self-Attention Net (DSN) Since a user might explore multiple geologically clustered cities before travel, e.g., cities around the Greater Boston area or the Greater New York area, we design a distance aware self-attention net that represents the k-th city in $\mathcal{J}_i$ as

$$\hat{\mathbf{c}}_{ik}^{on} = \sum_j \left(1 - \frac{exp(dist(c_{ik}^{on}, c_{ij}^{on}))}{\sum_t exp(dist(c_{ik}^{on}, c_{it}^{on}))}\right) \mathbf{c}_{ij}^{on}, \tag{1}$$

where $dist(,)$ is the Haversine distance between two city centers and $\mathbf{c}_{ij}^{on}$ is the city embedding through the embedding layer shared with the offline module.

Target-aware Attention Net (TAN) This attention net serves to extract the information related to the actual visited city c_i^{off} from the online behavior sequence, and represent it as

$$\mathbf{c}_{\mathcal{J}_i} = \sum_j \frac{exp(\langle \mathbf{c}_i^{off}, \mathbf{c}_{ij}^{on} \rangle)}{\sum_t exp(\langle \mathbf{c}_i^{off}, \mathbf{c}_{it}^{on} \rangle)} \mathbf{c}_{ij}^{on}. \tag{2}$$

Information Gain Net (IGN) Intuitively, a user is less likely to visit a city if she has visited it before, and thus we should emphasize more on the her interested but unvisited cities. For this purpose, the information gain net subtracts $\mathbf{c}_{\mathcal{J}_i}$ (the output of **TAN**) from the output of **DSN** and obtains $\mathbf{h}_i^{on} = \sum_k (\hat{\mathbf{c}}_{ik}^{on} - \mathbf{c}_{\mathcal{J}_i})$ as the representation of the online to offline information gain.

4.4 Final Prediction Layer

For each of the candidate city, it would also pass through **DSN** and **TAN** and get $\hat{\mathbf{c}}_{Tk}^{on}$ and $\mathbf{c}_{\mathcal{J}_T}$. However, instead of subtracting them, we add them together and output $\mathbf{h}_{cand}^{on} = \sum_k(\hat{\mathbf{c}}_{Tk}^{on} + \mathbf{c}_{\mathcal{J}_T})$. For the final prediction, we concatenate $\mathbf{h}_i^{on}$ from each time step, $\mathbf{h}_{cand}^{on}$, and $\mathbf{h}_T$ from the offline module and feed them to a MLP layer followed by a sigmoid activation unit to obtain the final predicted score for the candidate city. The overall model is trained to minimize the cross entropy loss.

5 EXPERIMENTS

5.1 Experimental Setup

5.1.1 Dataset. As far as we know, there is no public dataset that contains both user's online behavior and offline mobility data. Filling this gap, we collected user authorized behavior and mobility data for the past three years from a commercial OTP. The collected dataset contains 1.76M users and 341 cities. We label the cities the user visited as positive samples, while the negative samples are randomly sampled from all other cities.

5.1.2 Metrics. We adopt the widely used standard metrics Recall@K and Precision@K for evaluation. For each test sample, both the recall and precision are calculated. The final reported results are averaged over all samples.

5.1.3 Comparison Methods. The eight baselines used in the experiments are the state-of-the-art algorithms for sequential recommendation and next POI recommendation respectively. Specifically, the sequential recommendation algorithms are **1**) GRU4Rec [6], **2**) SASRec [15], **3**) DSIN[3] and **4**) SINS [15]. The comparing next POI recommdenation algorithms include: **5**) STRNN [10], **6**) STGN [17], **7**) GeoSAN [9] and **8**) STAN [11].

5.2 Experimental Results

The comparison results of the proposed OOPIN model and the baselines are summarized in Table 1. Clearly, OOPIN consistently outperforms all of the baselines on both evaluation metrics in general. We have the following main observations: 1) the best method for sequential recommendations using the online behavior data, i.e., SINS, is better than the best method for next POI recommendations using the offline mobility data, which shows that the online exploration behavior, if properly modeled, is more predictive of the next destination. This also matches our observation in the ablation studies. 2) Among the methods for next POI recommendation, STAN achieves the best performance since it explicitly considers the spatial-temporal preferences with a POI graph. 3) OOPIN outperforms all the baselines by jointly modeling the spatial-temporal periodicity from the offline mobility data and the information gain from the online behavior data.

5.3 Ablation Studies

5.3.1 Effectiveness of offline and online modules. To evaluate the impact of the two key components, i.e., the offline mobility pattern extractor and the online multi-interests exploration module, we remove either respectively. From Table 2, we can observe a significant performance drop when either is removed. Moreover, removing the online module has larger impacts, which verifies the importance of leveraging user online behavior data.

Table 1: Performance comparison of different methods.

Methods	Recall@10	Precision@10	Recall@5	Precision@5
GRU4Rec	0.0841	0.0997	0.0532	0.0581
SASRec	0.1167	0.1336	0.0696	0.0753
DSIN	0.1275	0.1480	0.0727	0.0778
SINS	0.1438	0.1652	0.0814	0.0875
ST-RNN	0.1059	0.1068	0.0528	0.0537
STGN	0.1186	0.1191	0.0583	0.0585
GeoSAN	0.1254	0.1228	0.0627	0.0617
STAN	0.1293	0.1289	0.0684	0.0689
OOPIN	**0.1693**	**0.1948**	**0.0971**	**0.1064**

5.3.2 Effectiveness of several key ingredients. Furthermore, we remove or replace several key components with some conventional modules, to obtain several new variants. For example, we replace **DSN** with conventional self-attention without distances, w/o **TAN** means passing $\mathbf{c}_k^{off}$ directly to **IGN**, w/o **IGN** means replacing it with MLP, and w/o PGRU means adopting regular GRU without the periodicity signal. We observe that after removing **IGN**, both metrics drop significantly, which indicates the importance of information gain from the candidate destinations that have not been visited by the user. Furthermore, after replacing PGRU with regular GRU, the performances also decrease a lot, indicating the benefit of the spatial-temporal periodicity in the sequential dependencies.

Table 2: Ablation studies of the key ingredients of the model.

Methods	Recall@10	RelaImp	Precision@10	RelaImp
OOPIN w/o Online	0.1135	-32.96%	0.1379	-29.21%
OOPIN w/o DSN	0.1592	-5.97%	0.1851	-4.98%
OOPIN w/o TAN	0.1583	-6.50%	0.1827	-6.21%
OOPIN w/o IGN	0.1527	-9.81%	0.1769	-9.19%
OOPIN w/o Offline	0.1532	-9.51%	0.1813	-6.93%
OOPIN w/o PGRU	0.1578	-6.79%	0.1864	-4.31%
OOPIN (standard)	**0.1693**	0.0%	**0.1948**	0.0%

5.4 Online A/B Test

We also conduct online A/B test by deploying the OOPIN model to the "guess you like" interface on the corresponding OTP App for ten days in July 2021. The treatment group assigns higher scores to OOPIN predicted destinations on top of the original ranking scores. The control group still uses the original ranking scores. We use the ***click-through rate*** (*CTR* for short) as an indicator to evaluate the performance of the experiment. With OOPIN, the treatment group brings 3.73% CTR gain over the control group.

6 CONCLUSION

We address the travel destination prediction problem facing OTPs in this paper. We propose an online-offline periodicity-aware information gain network, OOPIN, which consists of an offline mobility pattern extractors, that models the sequential dependencies with periodicity awareness from the offline mobility data, and an online multi-interests exploration module that models the information gain from the destinations that the user interested in but not yet visited. Extensive experiments conducted both offline and online demonstrate the effectiveness of our proposed model.

REFERENCES

[1] Chen Cheng, Haiqin Yang, Michael R. Lyu, and Irwin King. 2013. Where You like to Go next: Successive Point-of-Interest Recommendation. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (Beijing, China) *(IJCAI '13)*. AAAI Press, 2605–2611.

[2] Jie Feng, Yong Li, Chao Zhang, Funing Sun, Fanchao Meng, Ang Guo, and Depeng Jin. 2018. DeepMove: Predicting Human Mobility with Attentional Recurrent Networks. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1459–1468. https://doi.org/10.1145/3178876.3186058

[3] Yufei Feng, Fuyu Lv, Weichen Shen, Menghan Wang, Fei Sun, Yu Zhu, and Keping Yang. 2019. Deep session interest network for click-through rate prediction. *arXiv preprint arXiv:1905.06482* (2019).

[4] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. 2012. Next place prediction using mobility markov chains. In *Proceedings of the first workshop on measurement, privacy, and mobility*. 1–6.

[5] Long Guo, Lifeng Hua, Rongfei Jia, Binqiang Zhao, Xiaobo Wang, and Bin Cui. 2019. Buying or Browsing?: Predicting Real-Time Purchasing Intent Using Attention-Based Deep Network with Multiple Behavior. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery amp; Data Mining* (Anchorage, AK, USA) *(KDD '19)*. Association for Computing Machinery, New York, NY, USA, 1984–1992. https://doi.org/10.1145/3292500.3330670

[6] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).

[7] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.

[8] Jina Kim, HongGeun Ji, Soyoung Oh, Syjung Hwang, Eunil Park, and Angel P. del Pobil. 2021. A deep hybrid learning model for customer repurchase behavior. *Journal of Retailing and Consumer Services* 59 (2021), 102381. https://doi.org/10.1016/j.jretconser.2020.102381

[9] Defu Lian, Yongji Wu, Yong Ge, Xing Xie, and Enhong Chen. 2020. Geography-aware sequential location recommendation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 2009–2019.

[10] Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. Predicting the next location: A recurrent model with spatial and temporal contexts. In *Thirtieth AAAI conference on artificial intelligence*.

[11] Yingtao Luo, Qiang Liu, and Zhaocheng Liu. 2021. Stan: Spatio-temporal attention network for next location recommendation. In *Proceedings of the Web Conference 2021*. 2177–2185.

[12] Qijie Shen, Wanjie Tao, Jing Zhang, Hong Wen, Zulong Chen, and Quan Lu. 2021. SAR-Net: A Scenario-Aware Ranking Network for Personalized Fair Recommendation in Hundreds of Travel Scenarios. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4094–4103.

[13] Qijie Shen, Hong Wen, Wanjie Tao, Jing Zhang, Fuyu Lv, Zulong Chen, and Zhao Li. 2022. Deep Interest Highlight Network for Click-Through Rate Prediction in Trigger-Induced Recommendation. *arXiv preprint arXiv:2202.08959* (2022).

[14] Ke Sun, Tieyun Qian, Tong Chen, Yile Liang, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2020. Where to go next: Modeling long-and short-term user preferences for point-of-interest recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 214–221.

[15] Qiaoyu Tan, Jianwei Zhang, Jiangchao Yao, Ninghao Liu, Jingren Zhou, Hongxia Yang, and Xia Hu. 2021. Sparse-interest network for sequential recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 598–606.

[16] Wanjie Tao, Yu Li, Liangyue Li, Zulong Chen, Hong Wen, Peilin Chen, Tingting Liang, and Quan Lu. 2022. SMINet: State-Aware Multi-Aspect Interests Representation Network for Cold-Start Users Recommendation. In *AAAI*.

[17] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Fuzhen Zhuang, Jiajie Xu, Zhixu Li, Victor S Sheng, and Xiaofang Zhou. 2020. Where to go next: A spatio-temporal gated network for next poi recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2020).