

Where to Move Next: Zero-shot Generalization of LLMs for Next POI Recommendation

Shanshan Feng^{1,2†}, Haoming Lyu^{3†}, Fan Li^{4*}, Zhu Sun^{1,2}, Caishun Chen^{1,2}

¹Centre for Frontier AI Research, A*STAR, Singapore;

²Institute of High Performance Computing, A*STAR, Singapore;

³Nanyang Technological University, Singapore;

⁴Hong Kong Polytechnic University, Hong Kong, China

victor_fengss@foxmail.com; haominglyu@gmail.com; fan-5.li@polyu.edu.hk;
sunzhuntu@gmail.com; chen_caishun@cfar.a-star.edu.sg

Abstract—Next Point-of-interest (POI) recommendation provides valuable suggestions for users to explore their surrounding environment. Existing studies rely on building recommendation models from large-scale users’ check-in data, which is task-specific and needs extensive computational resources. Recently, the pretrained large language models (LLMs) have achieved significant advancements in various NLP tasks and have also been investigated for recommendation scenarios. However, the generalization abilities of LLMs still are unexplored to address the next POI recommendations, where users’ geographical movement patterns should be extracted. Although there are studies that leverage LLMs for next-item recommendations, they fail to consider the geographical influence and sequential transitions. Hence, they cannot effectively solve the next POI recommendation task. To this end, we design novel prompting strategies and conduct empirical studies to assess the capability of LLMs, e.g., ChatGPT, for predicting a user’s next check-in. Specifically, we consider several essential factors in human movement behaviors, including user geographical preference, spatial distance, and sequential transitions, and formulate the recommendation task as a ranking problem. Through extensive experiments on two widely used real-world datasets, we derive several key findings. Empirical evaluations demonstrate that LLMs have promising zero-shot recommendation abilities and can provide accurate and reasonable predictions. We also reveal that LLMs cannot accurately comprehend geographical context information and are sensitive to the order of presentation of candidate POIs, which shows the limitations of LLMs and necessitates further research on robust human mobility reasoning mechanisms.

Index Terms—LLMs, Next POI Recommendation, Zero-shot, Spatial-Temporal Data

I. INTRODUCTION

Recent years have witnessed the rapid development of location-based social networks (LBSNs) such as Foursquare and Facebook Places, where users can share their geographical positions by checking in points of interest (POI) on social networks. POIs usually denote the specific geographical locations that some users might find useful or interesting, such as coffee shops and libraries. Based on the check-in records, we can learn the user’s mobility movement patterns and further recommend appropriate POIs for users to visit. The POI recommendation task [1] is of great value in real-world scenarios, as it can help users to better explore their

surroundings, attract potential consumers for business holders, and increase the revenue of service platforms.

Compared with conventional POI recommendation tasks, the next POI recommendation task [2] focuses specifically on predicting the user’s next likely visit, which is more challenging. As presented in Figure 1, given a user’s check-in trajectory $\{l_1, l_2, l_3, l_4\}$, it aims to recommend the next location to visit. The next POI recommendation problem has attracted extensive research interest and various recommendation models have been developed [3]–[8]. However, existing methods need construct and train recommendation models from large-scale users’ check-in data, which consumes extensive computational resources. In addition, these recommendation methods are task-specific and lack generalization capability. In this work, we examine the next POI recommendation task from another perspective. Instead of training a task-specific recommendation model, we attempt to leverage the general-purpose pretrained large language models for generating sequential POI suggestions, which has not been explored before.



Fig. 1. An example of the next POI recommendation task. Each POI is associated with geographical coordinates and category information. Given the recent trajectory, e.g., $\{l_1, l_2, l_3, l_4\}$, it aims to predict POIs to visit subsequently.

LLMs not only have achieved remarkable results in various natural language processing tasks, but also have shown impressive performance in many domains. First, the geographical information can be extracted from the pretrained LLMs and further used for spatial-temporal studies. For example, Manvi *et al.* [9] prove that LLMs embed remarkable spatial information, and Gurnee *et al.* [10] find that LLMs learn

(†Both authors contributed equally to this research; *Corresponding author.)

linear representations of space and time across multiple scales. Roberts *et al.* [11] examine the degree to which GPT-4 acquires factual geographical knowledge and interpretative reasoning. Second, LLMs have been used for analyzing user mobility behaviors and spatial trajectory patterns [12]–[14]. However, these preliminary studies do not fully consider the geographical correlations and focus on different targets, e.g., anomaly trajectory detection and public event prediction, which are dissimilar to our work. Last, several LLM-based next-item methods [15]–[18] have been proposed and obtained promising zero-shot sequential recommendation performance on movie and E-commerce dataset, etc. Current solutions, unfortunately, fall short of capturing crucial aspects for next-Point of Interest (POI) recommendation tasks, specifically the geographical correlations and sequential transitions. Consequently, there is a pressing need to explore effective strategies for leveraging LLMs to address personalized user mobility recommendations.

This study investigates the application of pretrained LLMs in modeling human check-in trajectory data. Utilizing LLMs for this purpose presents a non-trivial challenge, as LLMs are inherently designed and optimized for language processing, making direct usage impractical for location prediction. To overcome this limitation, we introduce a novel framework named LLMmove, aiming to seamlessly integrate human movement prediction with language modeling. One of the key differences between next POI recommendation and the next-item recommendation [15]–[18] is the geographical correlations in user movements since users tend to visit close locations rather than far away places. This assumption is consistent with Tobler’s first law of geography, “Everything is related to everything else, but near things are more related than distant things”, which is the fundamental assumption used in spatial analysis. Specifically, we present the check-in data into *long-term check-ins* and *recent check-ins*, which reflects the user’s long-term spatial preference and the current spatial preference, respectively. For each POI in the *candidate set*, we calculate its geographical distance¹ from the user’s current position (indicated by the last check-in). Here, we incorporate four important factors for the next POI recommendation: long-term preference, current preference, geospatial distance, and potential sequential transitions. Then, by considering these requirements, LLMs are instructed to recommend Top-K POIs and provide explanations for the returned recommendations.

We conduct extensive experiments on two widely used real-world datasets for next POI recommender systems, which yield several significant insights. Our empirical evaluations showcase the promising zero-shot recommendation capabilities of LLMs, providing relatively accurate and reasonable predictions. However, it becomes evident that LLMs struggle with accurately grasping geographical context information and exhibit sensitivity to the order in which candidate POIs are presented. These limitations underscore the need for further

¹We also explore many strategies to use ChatGpt to directly compute the spatial distances based on the POIs’ coordinates, but cannot obtain accurate results. Hence, we calculate the distances and utilize them as input.

research to develop robust human mobility reasoning mechanisms in conjunction with LLMs.

The contributions of this work are summarized as follows:

- We investigate a novel research task, which explores the zero-shot generalization of LLMs to address the next POI recommendation. To the best of our knowledge, this is the first work to utilize the LLMs for POI recommendations.
- We develop a novel prompting framework, namely LLM-move, to incorporate various factors for sequential POI recommendation, including user spatial preferences, geographical distances, and sequential transitions.
- We conduct extensive experiments on two real-world datasets and derive several findings. The empirical results demonstrate the effectiveness of the proposed framework. The datasets and codes are available at <https://github.com/LLMMove/LLMMove>.

II. RELATED WORK

A. Next-POI Recommendation

As an important human mobility mining task, the next-POI recommendation problem captures the users’ complex personalized check-in behaviors, where various factors play essential roles including individual interests, continuous movement patterns, and spatial-temporal influence, etc. Recently, the next POI recommendation has attracted extensive research interests and a large variety of approaches have been developed [4]–[8], [19]–[21]. However, existing approaches require constructing and training recommendation models using extensive users’ check-in data, demanding significant computational resources. Moreover, these task-specific recommendation models fall short in providing zero-shot POI recommendations for users. Different from them, this work aims to generate the next POI suggestions without the need for task-specific training, a direction not previously explored.

B. LLM-Based Recommender Systems

Very recently, the LLMs have been exploited for the recommendation tasks [22]. Although LLMs are not specifically designed for capturing user-item interactions, their proficiency in understanding textual information and robust generative capabilities, including providing explanations and justifications, holds significant promise for improving recommendations. An illustrative example is the generative GPT4Rec framework proposed in [23], which treats the recommendation task as a query generation and searching procedure.

Several LLM-based recommendation methods address sequential recommendation problems. Harte *et al.* [24] propose three variants: LLM Embeddings, Fine-Tuned LLM, and LLM-enhanced Sequential Model. Wang *et al.* [15] introduce Zero-Shot Next-Item Recommendation with a prompting strategy guiding GPT-3 through user preferences, historical items, and top-K recommendations. Liu *et al.* [16] evaluate ChatGPT in five recommendation scenarios, employing zero-shot and few-shot prompt strategies for next-item prediction based on past sequential behaviors. [17] enhances sequential recommendations with a recency-focused prompting method.

Dai *et al.* [18] combine ChatGPT with information retrieval for improved recommendation capabilities. However, these prompt-based methods lack consideration for geographical information, hindering their effectiveness in solving the next POI recommendation task.

C. LLMs for User Mobility Patterns

The exploration of leveraging pre-trained models for modeling geographical spatial data has garnered increasing research attention. Two main approaches emerge: training geospatial pre-trained models and utilizing open-accessible Large Language Models (LLMs) for geospatial tasks. Mai *et al.* [25] and Balsebre *et al.* [26] focus on geospatial foundation models. Based on the open-sourced LLaMA model, Deng *et al.* [27] develop a foundation language model for understanding and utilizing geoscience knowledge. For open-accessible LLMs, studies like [11], [28], [29] explore tasks like population description, economic livelihood measurement, and route planning. [9] and [10] highlight that LLMs capture spatial information and acquire coherent knowledge about space and time.

Approaches like [12]–[14] employ pre-trained LLMs for various human mobility prediction tasks, such as anomaly detection using LLMs [13], predicting travel demand under public events [14]. It's worth mentioning that [12] introduces the LLM-Mob framework, utilizing accessible LLMs for learning mobility data. While it accounts for both long-term and short-term dependencies, its primary focus is on incorporating temporal information into human mobility sequences. However, the framework is specifically designed for the time-aware location prediction task, limiting recommendations to historically visited places. In essence, it does not provide recommendations for new locations and does not consider the geographical information of places.

III. PROBLEM STATEMENT

Let $L = \{l_1, l_2, \dots, l_{|L|}\}$ represent a set of Points of Interest (POIs), with each POI denoted as $\langle Id, Cat, Lat, Lon \rangle$. Here, Id denotes the unique ID for the respective POI, Cat indicates its category (e.g., Gym or Train station) providing semantic information, and Lat and Lon signify geographical coordinates, specifying latitude and longitude, respectively. Each check-in is represented as a tuple $c_{l,t}^u = \langle u, l, t \rangle$, indicating that user u visited POI l at timestamp t . A trajectory $traj = \{c_{l_1, t_1}^u, c_{l_2, t_2}^u, \dots, c_{l_k, t_k}^u\}$ represents a sequence of POIs visited by a user within a short time interval (e.g., 24 hours in this study). Focusing primarily on POI sequences in this work, the trajectory is denoted as $traj = \{l_1, l_2, \dots, l_k\}$ to avoid ambiguity.

Building upon prior research [5], [8], the objective of next POI recommendation is to furnish a list of potential POIs that a user is likely to visit subsequently. Formally, given the historical check-ins of a specific user and their current trajectory $traj = \{l_1, l_2, \dots, l_k\}$, the aim is to predict the probable next POI l_{k+1} that will be visited in the near future.

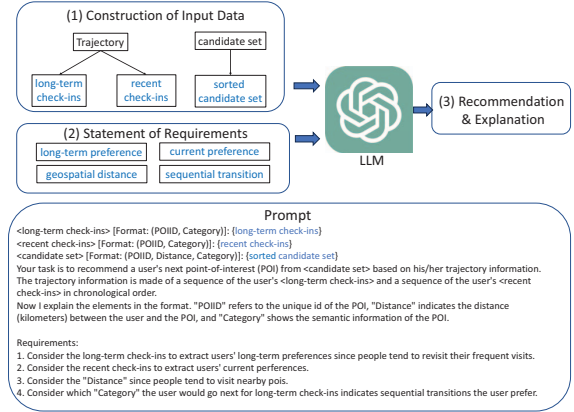


Fig. 2. The workflow of the designed LLMmove framework and the corresponding prompts.

IV. METHODOLOGY

Our study empowers Large Language Models (LLMs) for zero-shot next Point of Interest (POI) recommendations using a multi-step prompting strategy, namely LLMmove, as depicted in Figure 3. This LLMmove framework guides LLMs through three phases. Initially, it preprocesses data and incorporates background knowledge as input. Subsequently, it instructs the LLM to weigh four key factors: long-term and current user preferences, geographical distance, and sequential transitions. Finally, the LLM generates top- k POI recommendations with accompanying explanations.

A. Construction of Input Data

For personalized next POI recommendations, we leverage two types of user trajectory data: long-term check-ins capturing broader preferences and recent check-ins reflecting immediate interests. Additionally, we factor in the geographical distance of candidate POIs from the user's current location, recognizing its impact on travel behavior. These elements feed into the LLM, enabling it to suggest relevant and timely locations for each user.

B. Statement of Requirements

When analyzing user mobility behaviors, four key factors demand consideration. The first is long-term check-ins, offering insights into user preferences. The second, recent check-ins, mirror current contextual preferences. Thirdly, distance plays a role, as users tend to favor nearby Points of Interest (POI). Additionally, we delve into sequential transition patterns, exploring the flow between continuous categories in users' long-term check-ins.

C. Recommendation and Explanation

In this stage, we guide the LLM to generate the next POI recommendation along with reasons for the suggestion. Leveraging background knowledge, including both long-term and recent check-ins, and considering the candidate side, the LLM is instructed to incorporate the four requirements into the user

movement suggestion. Given the conditions and requirements, the LLM produces recommendations and explanations.

V. EXPERIMENTS

A. Experimental Setup

1) *Datasets*: To assess the efficacy of leveraging LLMs for next-POI recommendation, we employ two widely-used datasets: NYC [30] and TKY [30], following the experimental setting of [8]². Both the NYC and TKY datasets were collected from the Foursquare from April 2012 to February 2013, constituting location-based social network check-in records. NYC denotes the check-in data in New York City, while TKY covers the check-in data from Tokyo. Each check-in record includes the user, POI, POI category, geographical coordinates, and timestamp. The check-in records are organized in chronological order, with the first 80% serving as the training set, the subsequent 10% as the validation set, and the last 10% as the test set. The statistics of check-in datasets are reported in Table I. Here, #Test-Traj denotes the number of trajectories in test cases. For a given test trajectory $Traj = \{l_1, l_2, \dots, l_k\}$, we utilize the first $k - 1$ check-ins as the recent context, and the last visited POI l_k is considered as the ground truth POI. In each test case, the candidate set comprises the ground truth POI and 100 randomly sampled POIs, following a common practice in recommendation evaluation [31].

It’s noteworthy that historical records in the training set are exclusively used to reflect users’ long-term preferences. Instead of constructing sequential POI recommendation models, as in existing studies, we directly employ LLMs to perform zero-shot POI recommendations for users’ next moves.

TABLE I
STATISTICS OF TWO REAL-WORLD CHECK-IN DATASETS

Dataset	#Users	#POIs	#Category	#Checkins	#Test-Traj
NYC	1,048	4,981	318	103,941	1,364
TKY	2,282	7,833	290	405,000	4,610

2) *Evaluated Methods*: We assess the performance of the following approaches:

- **Popu**: This method selects the most popular POIs, i.e., the frequently visited locations by users.
- **Dist**: This approach directly chooses the nearest locations, i.e., the locations with the shortest distance.
- **CZSR** [16]: It develops a set of prompts for different recommendation scenarios, and we choose the zero-shot sequential recommendation prompt as the compared baseline in this work.
- **LLMRank** [17]: LLMRank regards the recommendation task as a ranking task, where historical interactions serve as conditions, and LLMs are instructed to rank a set of candidates. We choose the recency-focused prompting to incorporate recent check-in information.
- **ListRank** [18]: It boosts the LLM’s recommendation capabilities by ranking policies. We use the list-wise ranking policy as the compared recommendation baseline.

²<https://github.com/ant-research/Spatio-Temporal-Hypergraph-Model>

- **LLMMob** [12]: It leverages LLMs to analyze human mobility data by considering both long-term and short-term dependencies. By incorporating the temporal information, it aims at solving the time-aware mobility prediction.
- **LLMMob(-Time)** [12]: To make a fair comparison with other baselines, it removes the temporal information in the LLMMob framework.
- **LLMMob(+Geo)**: It extends the LLMMob approach [12] by additionally considering geographical influence, which is the same with the setting of LLMmove.
- **LLMmove**: LLMmove extracts user preferences, geographical influence, and sequential transitions, which constitute our proposed prompting strategy.

For all the evaluated LLM-based methods, we use the *gpt-3.5-turbo* as the default LLM for a fair comparison.

3) *Evaluation Metrics*: Consistent with prior studies on POI recommendation [5], [8], we employ two widely used performance evaluation metrics: Top-K accuracy rates (Acc@K) and Mean Reciprocal Rank (MRR). Acc@k assesses whether the ground truth POI is present in the Top-K recommended list, while MRR considers the ranking position of the ground truth in the sorted recommended list. Given the n test cases, Acc@k and MRR are defined as:

$$Acc@k = \frac{1}{n} \sum_1^n hits(rank_{gt} \leq k), MRR = \frac{1}{n} \sum_1^n \frac{1}{rank_{gt}}. \quad (1)$$

The function $hits()$ is an indicator: it returns 1 if the condition is true, and 0 otherwise. $rank_{gt}$ represents the rank of the ground truth next POI in the recommended list.

B. Next-POI Recommendation Results

The empirical results for next-POI recommendation are presented in Table II. The observations are as follows: (1) Popu demonstrates satisfactory performance, suggesting users’ inclination towards popular places. (2) Dist shows promising results, surpassing existing sequential zero-shot recommendation methods significantly. This underscores the crucial role of geographical distance in the next POI recommendations. (3) Applying existing sequential zero-shot methods (CSR, LLMRank, ListRank) for predicting users’ next movements is unfeasible, as they lack geospatial consideration. (4) LLMMob achieves relatively higher scores on both datasets, indicating its effectiveness in modeling user mobility data. By comparing the three variants of LLMMob, we can learn that both temporal and spatial information are beneficial for predicting future movements. (5) The proposed LLMmove attains the best performance, highlighting the advantages of the prompting strategy. Different from other baselines, it can make full use of geographical distance and sequential transition patterns. In particular, LLMmove outperforms the LLMMob(+Geo) by effectively leveraging spatial distance and personalized sequential patterns in user check-in behaviors. Overall, the proposed LLMmove excels in the next POI recommendation, showcasing its efficacy in harnessing LLM capabilities and learning user movement trajectories.

TABLE II
THE EXPERIMENTAL RESULTS OF NEXT-POI RECOMMENDATION ON REAL-WORLD DATASETS.

Methods	NYC				TKY			
	Acc@1	Acc@5	Acc@10	MRR	Acc@1	Acc@5	Acc@10	MRR
Popu	0.0500	0.2200	0.2750	0.1168	0.2300	0.4000	0.5050	0.3148
Dist	0.3702	0.5700	0.6195	0.4452	0.2700	0.4850	0.5450	0.3682
CZSR	0.1600	0.2400	0.2800	0.1903	0.1300	0.1700	0.1800	0.1461
LLMRank	0.0400	0.1150	0.1750	0.0738	0.0100	0.0850	0.1600	0.0469
ListRank	0.1100	0.1700	0.2050	0.1347	0.1250	0.1500	0.1750	0.1357
LLMMob	0.3600	0.5450	0.6050	0.4384	0.3350	0.4800	0.5150	0.3873
LLMMOB(-time)	0.2750	0.5550	0.6500	0.3911	0.2750	0.4550	0.4900	0.3505
LLMMob(+Geo)	0.3850	0.5850	0.6550	0.4703	0.3200	0.5500	0.6100	0.4091
LLMmove	0.5200	0.6100	0.6650	0.5585	0.4200	0.5800	0.6250	0.4847

TABLE III
THE ABLATION STUDIES OF LLMMOVE ON THE NYC DATASET.

Variants	Acc@1	Acc@5	Acc@10	MRR
LLMmove	0.5200	0.6100	0.6650	0.5585
-LP	0.4900	0.5800	0.6200	0.5266
-RP	0.5100	0.6100	0.6500	0.5521
-Geo	0.4800	0.5750	0.6250	0.5201
-Seq	0.5250	0.5850	0.6400	0.5554

TABLE IV
THE IMPACT OF THE ORDER OF CANDIDATE POI ON THE NYC DATASET.

Order	Acc@1	Acc@5	Acc@10	MRR
(a) Dist-asc	0.5200	0.6100	0.6650	0.5585
(b) Dist-des	0.2000	0.3000	0.3250	0.2398
(c) Rand	0.3250	0.4500	0.5150	0.3854
(d) Freq-asc	0.3600	0.4450	0.4950	0.4060
(e) Freq-des	0.4400	0.5400	0.6400	0.4900

C. Ablation Studies

We conduct ablation studies to examine the impacts of different factors: long-term preference (LP), recent preference (RP), geographical influence (Geo), and sequential transition (Seq). The empirical results are presented in Table III, yielding the following observations: (1) The considerable performance gap between LLMmove-LP and LLMmove indicates the importance of long-term preference. (2) LLMmove-RP shows results close to LLMmove, indicating that current preference might not be significantly influential in LLMs. Despite the acknowledged relevance of a user’s current interest in POI recommendations, the LLMs may struggle to fully utilize this factor in zero-shot scenarios, lacking collaborative information from other users. (3) LLMmove-Geo exhibits the lowest performance, underscoring the essential role of geographical distance in next POI recommendations. (4) LLMmove-Seq attains results similar to LLMmove. Although it shows slightly better results on Acc@1, its performance is inferior on Acc@5, Acc@10, and MRR. The relatively lower scores of Acc@5 and Acc@10 imply that, without explicitly considering this requirement, LLM may overlook potential POIs associated with sequential transitions.

Moreover, we made an interesting observation: the order of the candidate set significantly affects the recommendation performance. To explore this impact, we compared several ways to present the candidate POIs: (a) sorted by their distances

in ascending order (Dist-asc); (b) sorted by their distances in descending order (Dist-des); (c) random sort (Rand); (d) sorted by their category frequencies in ascending order (Freq-asc); (e) sorted by their category frequencies in descending order (Freq-des). The results are reported in Table IV. Notably, presenting POIs with higher probabilities (e.g., shorter distance or higher popularity) at the beginning of the candidate set leads to remarkable improvement. For example, the performance of Dist-asc is significantly better than Rand and Dist-des. This improvement could be attributed to the LLM’s limited capability to handle a large number of candidates. As an empirical choice, we adopt the Dist-asc ordering to present POI candidates in our work.

D. Case Studies

To develop an intuitive understanding of LLMmove’s reasoning ability and interpretability, we select and analyze some example results from the NYC dataset, as shown in Table V. Case 1 and case 2 demonstrate that LLMmove can deliver accurate recommendations accompanied by plausible reasons. These reasons align with the given requirements, showcasing the model’s ability to reflect the specified factors for recommendations. However, LLMmove fails to give reasonable recommendations sometimes. From Case 3, we can observe that LLMmove points out the user’s frequently visited POI ‘1823’ (Bar) twice, which may be attributed to limited categories in the data. Additionally, when faced with insufficient sequential transitions, LLM may struggle to make accurate predictions.

VI. CONCLUSION

In this study, we concentrate on harnessing the capabilities of Large Language Models (LLMs) for the zero-shot next POI recommendation task. Our approach considers both the users’ long-term preferences and current preferences, as well as geographic spatial distance and sequential transitions in user mobility behaviors. To integrate these factors, we introduce a novel prompt strategy aimed at generating top-K subsequent POI recommendations along with the rationale for suggestions. Extensive experiments conducted on two real-world datasets demonstrate that our proposed method significantly outperforms existing LLM-based baselines, showcasing its effectiveness in the next POI recommendation. However, our findings highlight potential challenges in spatial reasoning and understanding geographical information, emphasizing the need

TABLE V
CASES OF REASONS. THE RED TEXT HIGHLIGHTS THE REQUIREMENTS,
AND THE GROUND TRUTH POI IS MARKED IN BLUE.

<p>Case 1 (Correct)</p> <p>“recommendation”: [“4975”, “1449”, “3864”, “2713”, “2525”, “3513”, “356”, “2836”, “3245”, “4321”].</p> <p>“reason”: “Based on the user’s long-term check-ins, the user frequently visits POI ‘1449’ (Gym / Fitness Center) and ‘2836’ (Neighborhood). Therefore, it is likely that the user will visit these places again. Additionally, considering the recent check-ins, the user has visited POI ‘1449’ (Gym / Fitness Center) and ‘3864’ (Neighborhood) recently, indicating a current preference for these places. Furthermore, the user tends to visit nearby pois, so POI ‘4975’ (Home (private)) and ‘2713’ (Subway) are also recommended. Finally, considering the sequential transitions in the long-term check-ins, the user may visit POI ‘3245’ (Elementary School) after visiting ‘2836’ (Neighborhood).”</p> <p>“groundtruth”: 4975</p>
<p>Case 2 (Correct)</p> <p>“recommendation”: [“1395”, “1494”, “646”, “327”, “4853”, “2442”, “2932”, “2222”, “212”, “1043”].</p> <p>“reason”: “Based on the user’s long-term check-ins, the user frequently visits ‘Plaza’ (POIID: 1494) and ‘Performing Arts Venue’ (POIID: 1709). Therefore, the user is likely to prefer places in the ‘Plaza’ category. Additionally, considering the recent check-ins, the user’s most recent check-in was at a ‘Park’ (POIID: 2379). Taking into account the distance, the closest POI in the candidate set is ‘Law School’ (POIID: 1395). Therefore, the recommended POIs are the ten most probable places in the candidate set that match the user’s preferences and are closest in distance.”</p> <p>“groundtruth”: 1494</p>
<p>Case 3 (Incorrect)</p> <p>“recommendation”: [“4882”, “3252”, “1986”, “4604”, “1297”, “4676”, “3081”, “1122”, “1232”, “1304”].</p> <p>“reason”: “Based on the user’s long-term check-ins, the user frequently visits POIIDs ‘1823’ (Bar), ‘2279’ (Deli / Bodega), ‘3481’ (Home (private)) and ‘1823’ (Bar). Therefore, the user is likely to prefer places in the ‘Bar’, ‘Deli / Bodega’, and ‘Home (private)’ categories. Considering the recent check-ins, the user has visited POIIDs ‘242’ (Bookstore) and ‘1048’ (Toy / Game Store). Based on the distance, the user is likely to prefer nearby places. Therefore, the recommended POIIDs are the ten most probable places in the candidate set that match the user’s preferences and are nearby.”</p> <p>“groundtruth”: 356</p>

for future research to enhance the performance of generative POI recommendations.

REFERENCES

- [1] P. Sánchez and A. Bellogín, “Point-of-interest recommender systems based on location-based social networks: a survey from an experimental perspective,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 11s, pp. 1–37, 2022.
- [2] S. Feng, X. Li, Y. Zeng, G. Cong, Y. M. Chee, and Q. Yuan, “Personalized ranking metric embedding for next new poi recommendation,” in *IJCAI*, 2015, pp. 2069–2075.
- [3] S. Feng, G. Cong, B. An, and Y. M. Chee, “Poi2vec: Geographical latent representation for predicting future visitors,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [4] F. Yin, Y. Liu, Z. Shen, L. Chen, S. Shang, and P. Han, “Next poi recommendation with dynamic graph and explicit dependency,” in *AAAI*, vol. 37, no. 4, 2023, pp. 4827–4834.
- [5] S. Yang, J. Liu, and K. Zhao, “Getnext: trajectory flow map enhanced transformer for next poi recommendation,” in *SIGIR*, 2022.
- [6] S. Feng, L. V. Tran, G. Cong, L. Chen, J. Li, and F. Li, “Hme: A hyperbolic metric embedding approach for next-poi recommendation,” in *SIGIR*, 2020, pp. 1429–1438.
- [7] C. Duan, W. Fan, W. Zhou, H. Liu, and J. Wen, “Clsprec: Contrastive learning of long and short-term preferences for next poi recommendation,” in *CIKM*, 2023, pp. 473–482.
- [8] X. Yan, T. Song, Y. Jiao, J. He, J. Wang, R. Li, and W. Chu, “Spatio-temporal hypergraph learning for next poi recommendation,” in *SIGIR*, 2023, pp. 403–412.
- [9] R. Manvi, S. Khanna, G. Mai, M. Burke, D. Lobell, and S. Ermon, “Geollm: Extracting geospatial knowledge from large language models,” *arXiv preprint arXiv:2310.06213*, 2023.
- [10] W. Gurnee and M. Tegmark, “Language models represent space and time,” *arXiv preprint arXiv:2310.02207*, 2023.
- [11] J. Roberts, T. Lüddecke, S. Das, K. Han, and S. Albanie, “Gpt4geo: How a language model sees the world’s geography,” *arXiv preprint arXiv:2306.00020*, 2023.
- [12] X. Wang, M. Fang, Z. Zeng, and T. Cheng, “Where would i go next? large language models as human mobility predictors,” *arXiv preprint arXiv:2308.15197*, 2023.
- [13] Z. Zhang, H. Amiri, Z. Liu, A. Züfle, and L. Zhao, “Large language models for spatial trajectory patterns mining,” *arXiv preprint arXiv:2310.04942*, 2023.
- [14] Y. Liang, Y. Liu, X. Wang, and Z. Zhao, “Exploring large language models for human mobility prediction under public events,” *arXiv preprint arXiv:2311.17351*, 2023.
- [15] L. Wang and E.-P. Lim, “Zero-shot next-item recommendation using large pretrained language models,” *arXiv preprint arXiv:2304.03153*, 2023.
- [16] J. Liu, C. Liu, R. Lv, K. Zhou, and Y. Zhang, “Is chatgpt a good recommender? a preliminary study,” in *The 1st workshop on recommendation with generative models, CIKM*, 2023.
- [17] Y. Hou, J. Zhang, Z. Lin, H. Lu, R. Xie, J. McAuley, and W. X. Zhao, “Large language models are zero-shot rankers for recommender systems,” *arXiv preprint arXiv:2305.08845*, 2023.
- [18] S. Dai, N. Shao, H. Zhao, W. Yu, Z. Si, C. Xu, Z. Sun, X. Zhang, and J. Xu, “Uncovering chatgpt’s capabilities in recommender systems,” *Proceedings of the 17th ACM Conference on Recommender Systems*, 2023.
- [19] L. Zhang, Z. Sun, Z. Wu, J. Zhang, Y. S. Ong, and X. Qu, “Next point-of-interest recommendation with inferring multi-step future preferences,” in *IJCAI*, 2022, pp. 3751–3757.
- [20] J. Ou, H. Jin, X. Wang, H. Jiang, X. Wang, and C. Zhou, “Sta-tcn: Spatial-temporal attention over temporal convolutional network for next point-of-interest recommendation,” *TKDD*, vol. 17, no. 9, pp. 1–19, 2023.
- [21] Z. Sun, Y. Lei, L. Zhang, C. Li, Y.-S. Ong, and J. Zhang, “A multi-channel next poi recommendation framework with multi-granularity check-in signals,” *ACM Transactions on Information Systems*, vol. 42, no. 1, pp. 1–28, 2023.
- [22] W. Fan, Z. Zhao, J. Li, Y. Liu, X. Mei, Y. Wang, J. Tang, and Q. Li, “Recommender systems in the era of large language models (llms),” *arXiv preprint arXiv:2307.02046*, 2023.
- [23] J. Li, W. Zhang, T. Wang, G. Xiong, A. Lu, and G. Medioni, “Gpt4rec: A generative framework for personalized recommendation and user interests interpretation,” *arXiv preprint arXiv:2304.03879*, 2023.
- [24] J. Harte, W. Zörgdrager, P. Louridas, A. Katsifodimos, D. Jannach, and M. Fragkoulis, “Leveraging large language models for sequential recommendation,” in *Recsys*, 2023, pp. 1096–1102.
- [25] G. Mai, W. Huang, J. Sun, S. Song, D. Mishra, N. Liu, S. Gao, T. Liu, G. Cong, Y. Hu *et al.*, “On the opportunities and challenges of foundation models for geospatial artificial intelligence,” *arXiv preprint arXiv:2304.06798*, 2023.
- [26] P. Balsebre, W. Huang, G. Cong, and Y. Li, “Cityfm: City foundation models to solve urban challenges,” *arXiv preprint arXiv:2310.00583*, 2023.
- [27] C. Deng, T. Zhang, Z. He, Y. Xu, Q. Chen, Y. Shi, L. Fu, W. Zhang, X. Wang, C. Zhou, Z. Lin, and J. He, “K2: A foundation language model for geoscience knowledge understanding and utilization,” in *WSDM*, 2024.
- [28] Y. Zhang, C. Wei, S. Wu, Z. He, and W. Yu, “Geogpt: Understanding and processing geospatial tasks through an autonomous gpt,” *arXiv preprint arXiv:2307.07930*, 2023.
- [29] Z. Li and H. Ning, “Autonomous gis: the next-generation ai-powered gis,” *arXiv preprint arXiv:2305.06453*, 2023.
- [30] D. Yang, D. Zhang, V. W. Zheng, and Z. Yu, “Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 1, pp. 129–142, 2014.
- [31] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, “Neural collaborative filtering,” in *WWW*, 2017, pp. 173–182.