Context Trails: A dataset to study contextual and route recommendation

Pablo Sánchez*
psperez@icai.comillas.edu
Instituto de Investigación Tecnológica (IIT), Universidad Pontificia Comillas
Madrid, Spain

Alejandro Bellogín*
alejandro.bellogin@uam.es
Escuela Politécnica Superior
Universidad Autónoma de Madrid
Madrid, Spain

José L. Jorro-Aragoneses*
jose.jorro@uam.es
Escuela Politécnica Superior
Universidad Autónoma de Madrid
Madrid, Spain

Abstract

Recommender systems in the tourism domain are gaining increasing attention, yet the development of diverse recommendation tasks remains limited, largely due to the scarcity of public datasets. This paper introduces *Context Trails*, a novel dataset addressing this gap. *Context Trails* distinguishes itself by including not only user interactions with touristic venues, but also the itineraries (trails or routes) followed by users. Furthermore, it enriches existing item features (e.g., category, coordinates) with contextual attributes related to the interaction moment (e.g., weather) and the venue itself (e.g., opening hours). Beyond a detailed description of the dataset's characteristics, we evaluate the performance of several baseline algorithms across three distinct recommendation tasks: classical recommendation, route recommendation, and contextual recommendation. We believe this dataset will foster further research and development of advanced recommender systems within the tourism domain. Dataset is available at https://zenodo.org/records/15855966; further code available at https://github.com/pablosanchezp/ContextTrailsExperiments.

CCS Concepts

• **Information systems → Recommender systems**.

Keywords

Dataset, Context, Tourism, Route recommendation, Evaluation

ACM Reference Format:
Pablo Sánchez, Alejandro Bellogín, and José L. Jorro-Aragoneses. 2025. *Context Trails*: A dataset to study contextual and route recommendation. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys '25), September 22–26, 2025, Prague, Czech Republic*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Recommender systems are becoming increasingly present in our daily lives, helping users discover relevant content in different domains, such as e-commerce, entertainment, and, notably, tourism. In particular, Point-of-Interest (POI) recommendation plays a crucial role in enhancing user experience by suggesting relevant locations to visit when the users are exploring a city. However, unlike traditional domains like movies or books, POI recommendation is heavily influenced by contextual factors, such as time of day, weather, or the user's travel intention [27, 42].

While researchers can access a vast number of datasets for recommender systems – such as MovieLens (for movies), Amazon (for product reviews), or others[1] – there are few datasets specifically designed for POI recommendation that include rich contextual information (see Table 1). Public and well-known POI datasets, such as those from Location-Based Social Networks (LBSNs) like Foursquare[2] or Yelp[3], provide valuable data for POI recommendation, including basic POI metadata (e.g., categories or names) and user interactions (called *check-ins*) with timestamps. However, despite their usefulness, these datasets still miss stronger contextual signals such as weather or opening hours that are crucial for building realistic Context-Aware Recommender Systems (CARS).

To address this gap, we propose an extension of an existing Foursquare dataset, enriching it with additional contextual features (or providing the necessary code for that matter) such as POI categories, opening hours, and weather conditions at the time of each check-in. Herein, we present three sets of experiments using our dataset: traditional POI recommendation, route recommendation, and contextual recommendation, analyzing the behavior of a set of classical algorithms in the dataset. Thanks to this analysis, we consider that this extended dataset might be useful to the development of context-aware recommenders in both POI and route recommendation.

Hence, the contributions of this work are summarized as:

- We collect and provide the code to construct a new dataset covering user interactions as trails for three cities (New York, Petaling Jaya, and Tokyo). We enhance a previously public Foursquare dataset and integrate additional information about the POIs retrieved from the Foursquare API, along with weather conditions when check-ins occurred.
- We provide a comprehensive analysis of the dataset, including visualizations of POI spatial distribution, information on the user routes, and detailed statistics on weather conditions and the other collected contextual variables.
- We conduct benchmarking experiments across three recommendation tasks: classic Point-of-Interest, route, and context-aware recommendation.

*All authors contributed equally to this research.

[1] GitHub repository with links to datasets for recommendation tasks: https://github.com/ACMRecSys/recsys-datasets.
[2] Foursquare datasets: https://sites.google.com/site/yangdingqi/home/foursquare-dataset.
[3] Yelp dataset: https://business.yelp.com/data/resources/open-dataset/

RecSys '25, Prague, Czech Republic

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

2 Related work

This section presents the main recommendation tasks our dataset allows analyzing and a contextualization against other datasets.

2.1 Point-of-interest (POI) recommendation

The POI recommendation task consists of suggesting relevant venues (often referred as POIs) – such as restaurants, museums, or parks – users might want to visit, when traveling to a specific city [31]. These recommendations are typically derived by analyzing user preferences and historical behaviors, captured through check-ins registered in Location-Based Social Networks (LBSNs). These networks are online platforms where users voluntarily share their location by checking-in at specific venues, often accompanied by reviews or ratings [2, 39]. Some examples of these LBSNs include Foursquare, Gowalla, or Yelp. Check-ins are used to build user–POI interaction matrices, which are then exploited by recommendation models to infer preferences and suggest locations [9, 26].

Unlike traditional recommendation scenarios, such as movies or books, where user preferences tend to be relatively stable and not so dependent on temporal and geographic information, POI recommendation is inherently more dynamic and context-dependent. POI relevance often depends on situational factors such as the user's current location, the time of the day, weather conditions, or the user's intention (e.g., sightseeing or dining) [16, 36, 38]. To address these challenges, the field has evolved significantly in the last decade. Early approaches focused on Collaborative Filtering (CF) techniques, including similarity-based models and Matrix Factorization (MF), aimed at obtaining latent representations of users and venues [19, 44]. These techniques incorporated temporal dynamics, geographic information, and social influence to improve recommendation accuracy. More recently, deep learning has brought substantial advances, with architectures such as recurrent neural networks (RNNs) [14, 48], graph neural networks [45], and attention-based models [17, 43] which are applied to model sequential behavior, capture spatial correlations, and learn from heterogeneous sources.

2.2 Route recommendation

In addition to individual POI recommendation, a more challenging but highly relevant task is route recommendation, which involves suggesting an ordered sequence of POIs that a user may wish to visit within a single time frame. While traditional POI recommendation focuses on identifying a single relevant location based on user preferences and context, route recommendation must take into account multiple interdependent decisions, considering relationships between consecutive POIs, temporal constraints, and geographical proximity among other criteria [6, 21].

This task involves multi-objective optimization [8], where models must balance not only the relevance of each POI, but also factors like travel time, route coherence, time budgets, and starting and ending locations [13]. The combinatorial nature of the problem makes it particularly challenging: the number of potential routes increases exponentially with the number of candidate POIs, making efficient search, pruning strategies, and user personalization essential. Several techniques have been explored to address these challenges, ranging from heuristic-based proposals and graph-based models [46] to approaches that use sequence modeling (e.g., RNNs, transformers) and reinforcement learning to dynamically generate or rank route candidates [7, 47]. Some methods treat the problem as different sub-areas from the Operations Research field (e.g., orienteering and traveling salesman problems) with personalized and contextual constraints [12], while others rely on historical trajectory data to mine frequent patterns or transitions between POIs [18].

2.3 Context-aware recommendation

Recommender systems usually take into account only two variables: users and items. But a third dimension may frequently coexist, embedding the so-called "contextual information". Context is a multifaceted concept that has been used across different disciplines [20]. This additional information can be temporary, geographical, or any other type that can help the system to be more precise in its main goal; for example, information about the weather can help a tourist company to improve its recommendations. In this scenario, it has been shown that the incorporation of this type of information makes it possible to better characterize and profile users using more detailed data and, therefore, make better recommendations [37].

Starting with the seminal work of Adomavicius et al. [1], various authors have discussed the main challenges of CARS approaches; in that work, the authors illustrate the usage of this type of algorithms, focusing on travel and music domains, although foreseeing more interactive or adaptive scenarios like conversational applications. More recently, we find literature surveys specialized in particular domains regarding CARS methods, like social networks [35], cultural heritage [5], temporal aspects [4], or neural networks [20].

All these approaches can be characterized as belonging to one of the following three paradigms, depending on how context is incorporated into the recommendation model: i) *contextual pre-filtering*, where context is used to select the relevant set of data records and, then, recommendations are produced by using any classical (i.e., non-CARS) recommender system on the selected subset; ii) *contextual post-filtering*, where contextual information is ignored when training the algorithms, hence allowing to obtain recommendations again by using a classical RS trained with the whole data, subsequently, such recommendations are adjusted through a per-user contextualization by incorporating the contextual information - in other words, this mechanism refines the resultant recommendation list; and iii) *contextual modeling*, where context is directly integrated in the recommendation function [20].

2.4 Related datasets

Both the tourism domain and the context-aware recommendation area suffer from a lack of sufficiently large, diverse, and complete datasets. Based on a non-exhaustive search, but considering highly-cited related publications and recent surveys on these topics (i.e., [15, 20, 31]), we show in Tables 1 and 2 a collection of 15 CARS datasets and 8 POI and route datasets.

Regarding the datasets used in CARS approaches, besides standard statistics, we identify the domain and the contexts included in the data, and where several versions existed, we include a range of minimum-maximum values for that column. It should be noted that in some cases (as in those denoted as HC, from HyperCars [3]) these contexts are inferred from textual reviews, so the contextual information is not as certain as in other datasets. Also, in

Table 1: Datasets used in publications dealing with CARS approaches. K and M denote thousands and millions. Context abbreviations used: A (age), B (budget), C (companion), D (demographic), DS (driving style), G (goal), H (hunger level), L (location), LI (last interactions), M (mood), O (order), R (road type), RV (real/virtual), S (social), Sch (schedule), T (time), Ta (tag), TT (trip type), W (weather).

Dataset	Domain	Users	Items	Interactions	Contexts
Adom	Movie	0.1K	0.2K	1.5K	C, L
Comoda	Movie	0.1K	1.2K	2.3K	M, S, T, W
DePaulMovie	Movie	0.1K	0.1K	5K	C, L, T
In Car	Music	0.01K	0.2K	4K	DS, M, R, W
Food	Food	0.2K	0.01K	6.4K	H, RV
Foursquare	POI	0.2K-51K	0.3K-500K	0.5K-3.5M	D, L, T, W
Frappe	Apps	1K	4K	95K	L, T
HC-Gowalla	POI	24K	40K	1M	L, T, W
HC-Yelp	POI	312K	12.6K	1.1M	L, T, W
LastFM	Music	0.01K-3K	1.8K-174K	93K-19M	LI, O, Ta, T
MovieLens	Movie	0.7K-140K	1.6K-19K	31K-20M	A, T
STS	POI	0.3K	0.3K	2.5K	B, C, G, M, T, W
TripAdvisor	POI	1.2K-2.6K	1.5K-1.9K	4.7K-9.3K	TT
Weeplaces NY	POI	4.5K	16.1K	864K	W
Yelp	POI	5K-96K	13K-49K	144K-2.3M	LI, L, T
Context Trails	POI	85K	84K	1.3M	L, Sch, T, W

Table 2: Datasets used in publications dealing with POI and route recommendation approaches.

Dataset	Cities	Users	Items	Check-ins	Routes
Foursquare Global Scale	415	267K	3.7M	33.3M	NA
GeoLife	1	0.2K	≈17K	28M	17.6K
Gowalla	50	1.6K-107K	3.5K-1.3M	116K-6.4M	NA
Semantic Trails 2013	10K	256K	2.8M	18.6M	6.1M
Semantic Trails 2018	52K	400K	1.9M	11.9M	4M
Trip builder	3	22.6K	1.3K	133K	55.5K
VeronaCard	1	(unk)	0.1K	1.2M	250K
YFCC100M	1-7	0.9-6.5K	0.1K	17K-130K	4K-20K
Context Trails	3	85K	84K	1.3M	580K

other cases (like movies and music) debatable context features are used, like tags or age [33, 40, 49]. We also observe that POI is a popular domain where context has often been studied, because it offers richer context sources beyond time. Based on this sample, we conclude that our collected dataset is larger than most of these datasets, includes relevant contextual dimensions, and belongs to a domain that is interesting for the CARS community.

Regarding the datasets used in POI and route recommendation approaches, we also include the number of cities they cover and the routes, if available. It is important to note that many of these datasets contain a small number of POIs, which is surprising, especially considering the number of cities they claim to include. Since the field has devoted more effort on the POI, rather than the route, recommendation problem, it does not come as a surprise that not many datasets include some type of (explicit) route, trail, or trajectory of the users throughout the city. Our collected dataset is the one with more routes (except for Semantic Trails, which is, as we shall explain later, the data source we start from). There are also other datasets with more cities, but this is the first version of our dataset, and we are working on extending it with more cities, all of them with their corresponding routes and contexts.

2.5 Contextualizing our work

Our dataset offers the following features: i) user and item timestamped interactions, both in the form of check-ins, for traditional POI recommendation, and in the form of routes; ii) currently, a selection of popular but culturally diverse range of cities (New York, Petaling Jaya, and Tokyo); iii) together with temporal and geographical dimensions, which are frequently used in CARS literature [20], we include weather data (aligned to the timestamped interaction and the corresponding city) and schedules for the venues. While weather information is sometimes included in datasets of this domain (such as [36]), the schedules are frequently neglected, basically because the use case scenario of the simulated recommendations do not aim to be realistic (i.e., is the suggested venue really open when the user is expected to visit it?). Our dataset would allow answering this type of questions and create richer and realistic approaches.

3 The *Context Trails* dataset

3.1 Data collection

The dataset *Context Trails* we present here is based on three main data sources: SemanticTrails, VisualCrossing, and Foursquare:

- SemanticTrails: this dataset contains user routes (called trails by the authors) across various POIs in different cities around the world. It is available online[4] and was proposed in [23]. The dataset is divided into data from 2013 (check-ins extracted from the Global-Scale Check-in Dataset from footnote 2) and 2018. We focus on the 2018 version to include more recent interactions and avoid outdated information. The POIs visited by users are identified with Foursquare IDs. This dataset includes route and user identifiers, timestamps related to the check-ins, and categories and country of each POI. Note no geographical coordinates are included in this version of the dataset, even though this is a common feature among existing POI datasets (see Table 1).
- VisualCrossing[5] provides historical weather data, allowing to obtain information such as temperature, rainfall, and wind for different dates and locations around the world. The data was obtained with an hourly resolution.
- Foursquare Developers API[6] allows developers to obtain more specific details about the POIs, including their geographical coordinates and opening and closing hours.

There are multiple stages involved in building the *Context Trails* dataset. First, to promote more recent data and objectively defined user routes, we select the SemanticTrails dataset (2018 version). We choose all the trails from the cities of New York (NYC), Petaling Jaya (PTJ), and Tokyo (TOK), which correspond to Wikidata IDs Q60, Q864965, and Q308891, respectively. Second, we obtain the details of each POI in these cities using the Foursquare API: Foursquare ID, latitude, longitude, categories, price tier, average rating, total number of ratings, total number of user tips, and the specific opening and closing hours. Unfortunately, some POIs may be removed from the Foursquare API over time, making it impossible to retrieve any data for them (in particular, note our data collection corresponds

[4] See https://figshare.com/articles/dataset/Semantic_Trails_Datasets/7429076
[5] VisualCrossing: https://www.visualcrossing.com/
[6] Foursquare Developers: https://location.foursquare.com/developer/

Table 3: Statistics of the data collected for the three cities, including its Wikidata ID (WikID), number of users (|U|), POIs (|P|) and POIs with schedule (|P_s|), check-ins (|C|), and unique number of contexts as weather conditions (CT_w), temperature values (CT_t), and schedules (CT_s).

City	WikID	|U|	|P|	|C|	CT_w	CT_t	CT_s	|P_s|
NYC	Q60	1,649	1,461	4,849	10	430	53	1,177
PTJ	Q864965	18,346	18,618	153,543	4	122	153	17,207
TOK	Q308891	66,125	64,086	1,178,663	8	385	333	57,192

to November 2024). Specifically, in NYC, of a total of 1,461 POIs, this information could not be retrieved for 117 (8.0%). In Petaling Jaya, data was unavailable for 1,412 POIs of 18,618 (7.58%). Lastly, in TOK, no information was retrieved for 5,781 POIs out of a total of 64,086 (9.02%). Finally, we download the weather data from Visual Crossing, covering the overall period of the original Semantic Trails dataset (October 2017 to December 2018) [10]; once this data is obtained, we assign the corresponding weather information to each check-in according to the recorded timestamp. We report in Table 3 all this information for the three cities.

3.2 Data content

We provide the following data or the necessary code to retrieve it:

- **POI information:** For each POI, we provide code to retrieve its Foursquare ID, latitude, longitude, and categories (separated by "-"; each POI may belong to multiple categories). Additional attributes include price tier, average rating, total number of ratings, and total number of user tips. We also include a set of binary indicators that represent whether the POI is open during specific time intervals. The defined time windows are: *early morning* (00:00–06:00), *morning* (06:00–12:00), *afternoon* (12:00–18:00), and *night* (18:00–00:00). These availability indicators are computed separately for weekdays (Monday-Friday) and weekends (Saturday-Sunday). We use 1 to represent the POI is open at that interval, and 0 otherwise. As Foursquare defines over 1K categories, we also include the first-level category; this information is derived from the "Places Open Source & Pro/Premium Flat" file[7]. Specifically, we extract the first-level category by parsing the "category_label" field, which presents the full category hierarchy separated by the '>' symbol; the first segment corresponds to the first-level category.
- **Route & weather information:** For each route, we include the route ID, the user who followed the route, the associated Foursquare venue ID, and the timestamp of the check-in. In addition, we enrich each entry with contextual weather data at that specific time, including temperature (measured in Celsius), precipitation level (measured in mm), wind speed (measured in kph), type of precipitation, and sky conditions.
- **Training and test splits:** We provide the original training and test splits used in our experiments for both POI and route recommendation tasks. For POI recommendation, each record contains the user ID, venue ID, and timestamp. For route recommendation, we also include the route ID.

[7]"Places Open Source & Pro/Premium Flat" file: https://docs.foursquare.com/data-products/docs/categories

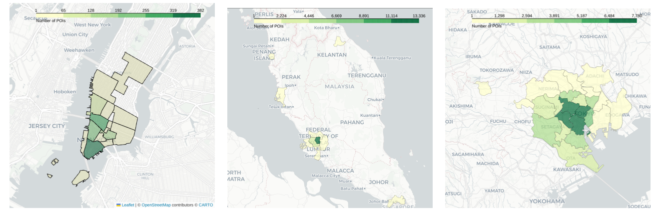

(a) New York City **(b) Petaling Jaya** **(c) Tokyo**

Figure 1: Choropleth map showing the distribution of points available in the three cities that conform *Context Trails*.

3.3 Data statistics

The main statistics for each city are reported in Table 3, where we specify the number of users, POIs, check-ins, contexts, and POIs with schedule. These statistics are computed for all cities, Tables 4 and 5 show the statistics when they are used for the recommendations problems considered herein.

We now describe the complete dataset for each city, prior to any split. Figure 1 shows the spatial distribution of POIs in each city, with darker green areas indicating the regions where POIs are more densely concentrated. For all cities, check-ins tend to be concentrated in the center, except for Petaling Jaya, which also contains check-ins from other areas. Figure 2 illustrates the distribution of route sizes, originally determined by the creators of the Semantic Trails dataset [23]. We observe that most of the routes are very short, in particular, less than 3 POIs.

Figure 3 presents the climographs for these cities, illustrating the monthly precipitation volumes and average temperatures from October 2017 to December 2018 (hourly resolution). To further analyze the impact of weather conditions on user activity, we also show in Figure 4 the number of check-ins for each weather condition in each city. Figure 5 presents the temperature distribution associated with the check-ins, while Figure 6 shows the distribution of different types of precipitation. Although there are notable differences between cities — for instance, New York exhibits a higher percentage of check-ins on partly cloudy days compared to Tokyo, and nearly all check-ins in Petaling Jaya occur on partly cloudy days — there are very few check-ins during snowy conditions. This lack of check-ins while snowing is due both to the unusual nature of snow in these locations and a general tendency for users to avoid checking in during extremely cold weather, as further evidenced by the temperature distribution in Figure 5 and the distribution of precipitation types in Figure 6.

Finally, Figure 7 shows the number of POIs in the most popular categories open during each time segment of the day. In particular, most POIs fall into the "dining and drinking" category, making it the most frequent in all time slots. Furthermore, the "retail" category exhibits increased activity after the early morning hours, indicating a temporal pattern in user engagement.

3.4 Dataset use

The *Context Trails* dataset can be used to train and evaluate traditional RS algorithms, such as CF or content-based [29]. However, its main strength lies in its contextual, POI, and route information;

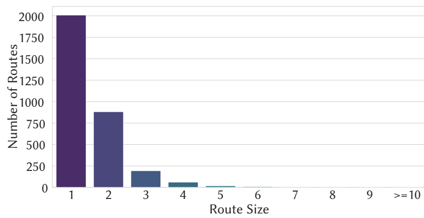

(a) New York City

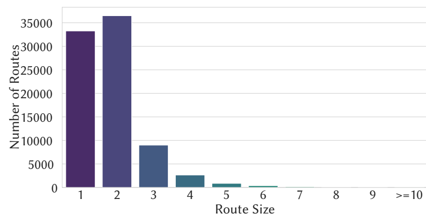

(b) Petaling Jaya

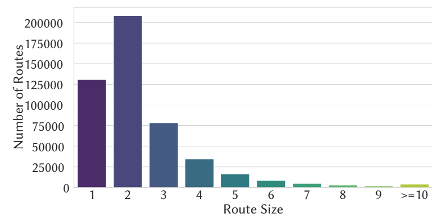

(c) Tokyo

Figure 2: Distribution of the route sizes.

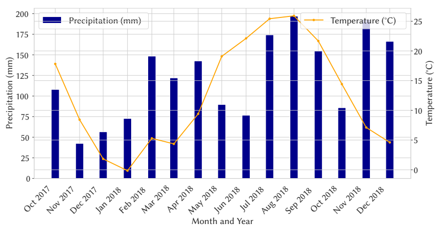

(a) New York City

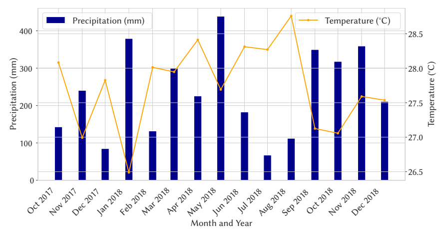

(b) Petaling Jaya

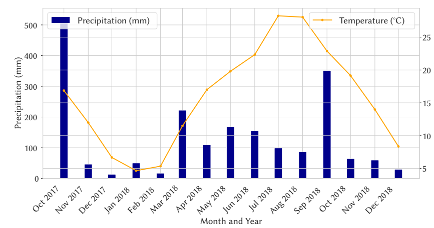

(c) Tokyo

Figure 3: Climograph of the three cities that conform *Context Trails* throughout the time covered in the data.

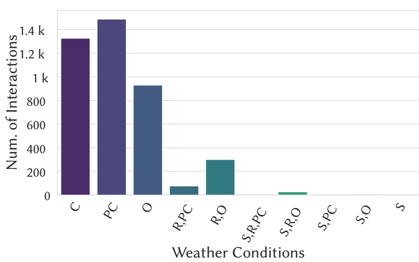

(a) New York City

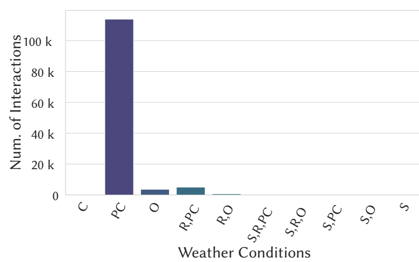

(b) Petaling Jaya

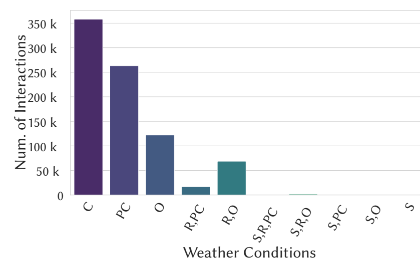

(c) Tokyo

Figure 4: Number of interactions for each weather condition included in the dataset. Conditions are: *C* (Clear), *PC* (Partially Cloudy), *O* (Overcast), *R* (Rain), and *S* (Snow). Combined conditions are represented by multiple abbreviations, e.g., *R,PC* stands for Rain and Partially Cloudy.

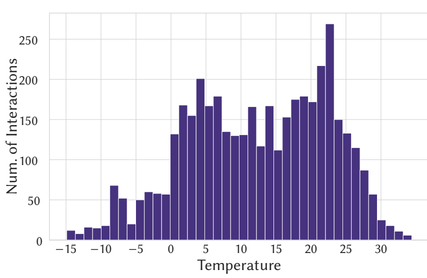

(a) New York City

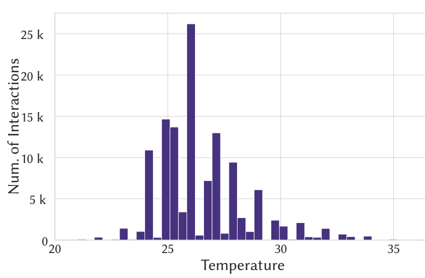

(b) Petaling Jaya

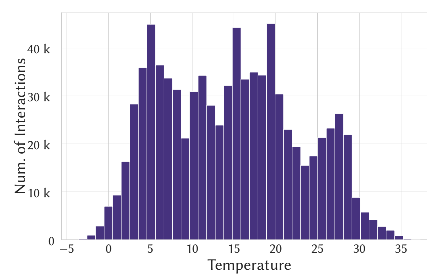

(c) Tokyo

Figure 5: Distribution of the temperatures.

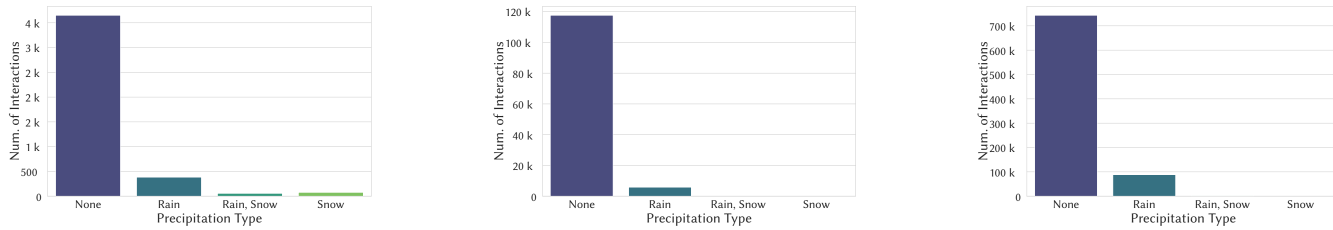

(a) New York City **(b) Petaling Jaya** **(c) Tokyo**

Figure 6: Distribution of the precipitation types.

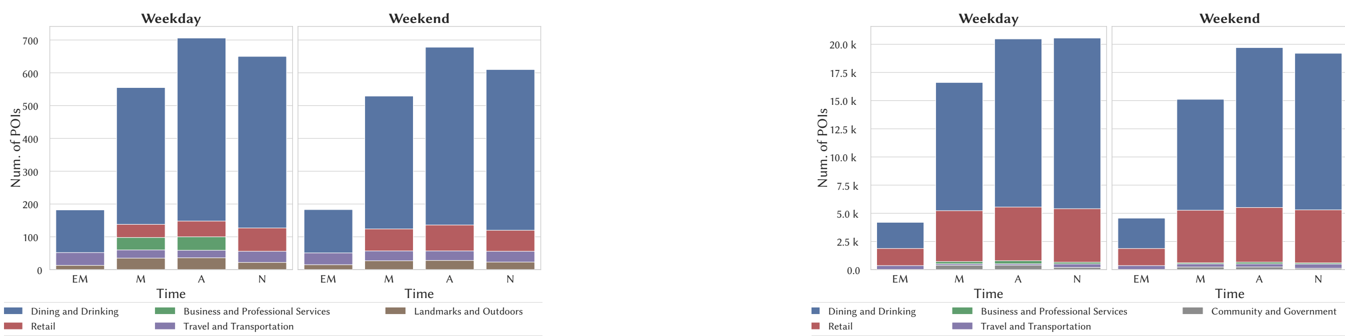

(a) New York City **(b) Tokyo**

Figure 7: Number of POIs open during each time segment of the day: Early Morning (EM), Morning (M), Afternoon (A), and Night (N). For NYC and TOK, two plots are shown displaying the number of open POIs per category, focusing on the top-10 most represented categories, in each plot corresponding to weekdays (left) or weekends (right).

hence, the nature of the data analysis and recommendation tasks it promotes is broad and varied. For example, it can be used to analyze how weather conditions influence human mobility patterns and POI visiting behaviors, providing information on user preferences in different atmospheric scenarios. It also supports the development and evaluation of adaptive models that adjust to external factors such as rain, temperature, or wind. In addition, it can also be used to generate weather-aware route recommendations, where suggested tours can avoid outdoor locations during unfavorable conditions.

In summary, the dataset provides valuable information to advance research at the intersection of human mobility, context-awareness, and urban computing.

4 Tasks and experiments

We now present the specific evaluation setup used for the addressed tasks: POI, route, and contextual recommendation. We detail algorithms, data partitioning strategies, and evaluation metrics.

4.1 Experimental setup

4.1.1 POI recommendation. In this task, we used the following traditional algorithms, tuned to maximize nDCG@5:

- Random (Rnd): perform recommendations randomly.
- Popularity (Pop): recommends to the target user the venues that have received the highest number of visits.
- UB and IB: non-normalized neighborhood CF approaches [24].
- BPR: Bayesian Personalized Ranking loss [28] used in MF.
- EASEr: Embarrassingly Shallow Autoencoders for Sparse Data proposed in [34].
- $RP^3\beta$: graph-based technique presented in [25].
- IRenMF: geographical weighted matrix factorization that models the geographical influence between neighbor POIs [19].
- GeoBPR: geographical BPR that incorporates geographical information by assuming that users will prefer to visit POIs close to the venues visited previously [44].
- H-PUM: a hybrid approach that combines three different recommenders: Pop, UB, and midpoint of the users.
- Skyline: perfect recommender that recommends the test set. This recommender serves as an upper bound for performance.

In this task, we first aggregate all user check-ins to ensure each user visited a given venue only once. We used the aggregated value to represent the user's level of interest in that specific venue, hence building a frequency user × POI matrix while retaining the timestamp of the user's most recent interaction with each venue. Next, we split the check-ins into training and test sets by applying a temporal partition. The oldest 80% of the dataset serves as the training set, while the remaining 20% is used for testing. Statistics in each processing step are shown in Table 4. As pointed out in Section 3.1, some information regarding the POIs could not be obtained. For those POIs, we considered their coordinates to be the center of the city, so that the recommenders that use geographical information could continue to generate recommendations for these POIs.

Table 4: Statistics of the three cities for POI recommendation configuration. For each city, we show the number of users |U|, POIs |P|, check-ins |C|, and density (δ).

City	Split	|U|	|P|	|C|	δ (%)
NYC	Original	1,649	1,461	4,849	0.20127
	Aggregated	1,649	1,461	4,161	0.17271
	Training	1,347	1,277	3,328	0.19337
	Test	397	459	833	0.45713
PTJ	Original	18,346	18,618	153,543	0.04495
	Aggregated	18,346	18,618	123,351	0.03611
	Training	16,316	16,102	98,680	0.03756
	Test	5,593	7,490	24,670	0.05889
TOK	Original	66,125	64,086	1,178,663	0.02781
	Aggregated	66,125	64,086	833,103	0.01966
	Training	59,361	58,012	666,482	0.01936
	Test	26,523	27,280	166,621	0.02303

4.1.2 Route recommendation. For this configuration, we used the following classical route recommenders, where every recommended route starts with the first POI visited by each test user:

- Baseline-R: recommends the first POI of the user in the test set, it is used to show the expected base performance.
- ClosestNN-R: route recommender suggesting the geographically nearest POI to the previously recommended one.
- MC-R: route recommender using a first-order Markov chain to suggest the next POI.
- FMC-R: same as the previous one, but considering the feature transition probability between the categories of the POIs.
- kNN-R: route recommender based on a k-nearest neighbors approach that computes the similarity between users using the Jaccard similarity. Candidate POIs are scored according to the similarity of the neighbor who visited them, accumulating scores from multiple neighbors if applicable.
- WG-R: route recommender that combines three components: distance-based weights (favoring geographically close POIs), transition-based weights (favoring frequently visited sequential POIs), and category-based weights (favoring POIs with similar category transitions).

Note that when computing the score for the next POI to recommend, the scores of multiple candidate POIs may be identical. In such cases, we resolve ties using one of two criteria: we select either the most popular one or the closest to the previously visited POI, depending on which one maximizes nDCG@5. Again, for those POIs for which no valid coordinates are available, the city center has been selected as the coordinates for these POIs.

For this configuration, we must consider that the routes should not be split. That is, every route must remain as a whole in the training or test sets. Based on this, we decided to split the data as follows: For each user with at least 2 different routes, the most recent route is assigned to the test set, provided it contains at least 4 check-ins; the remaining routes are allocated to the training set. No additional data filtering is applied. We show the statistics in each processing step for this task in Table 5.

4.1.3 Contextual recommendation. For the context-aware recommendation task, we selected 3 algorithms from those applied to the POI recommendation task (Section 4.1.1), and then applied a post-filter to retain only the items matching the target context [1]. The selected algorithms are: Random (C-Rnd), Popularity (C-Pop), and H-PUM (C-H-PUM).

Table 5: Statistics of the three cities for the route recommendation configuration. For each city, we show the number of users |U|, POIs |P|, check-ins |C|, density (δ), and trails |T|.

City	Split	|U|	|P|	|C|	δ (%)	|T|
NYC	Original	1,649	1,461	4,849	0.20127	3,186
	Training	1,649	1,446	4,763	0.19975	3,170
	Test	16	73	86	0.07363	16
PTJ	Original	18,346	18,618	153,543	0.04495	82,959
	Training	18,346	18,494	151,675	0.04470	82,568
	Test	390	1,060	1,868	0.45186	390
TOK	Original	66,125	64,086	1,178,663	0.02781	489,684
	Training	66,125	63,498	1,145,465	0.02728	483,814
	Test	5,870	8,461	33,198	0.06684	5,870

Each method was evaluated using three context configurations—time, weather, and their combination—employing the same data split as in Section 4.1.1, where the oldest 80% of interactions form the training set and the most recent 20% are used for testing. For each user–context pair in the test set, a POI ranking was generated based on the algorithm's strategy and contextual information (i.e., considering the original ranking from the algorithm and filtering according to the context(s)), and nDCG@5 was computed.

4.2 Results

In this section, we describe the results obtained in the POI recommendation task (Section 4.2.1), route recommendation (Section 4.2.2), and contextual recommendation (Section 4.2.3). Petaling Jaya is omitted for space constraints, but full results can be found in the online repository.

4.2.1 POI recommendation. We present in Table 6 the results obtained in the POI recommendation task. For each city, we present the performance of the recommenders in terms of ranking relevance (nDCG), novelty (EPC), diversity (Gini), and user coverage (UC), i.e., the number of users test to whom we are able to make recommendations. Based on the results, we observe that all recommenders (including the Skyline), except for Pop, Rnd, and H-PUM, do not achieve full coverage. This implies that there are test users for whom no recommendations are generated. This limitation arises because all generated recommendations are restricted to POIs present in the training set, and, because of the temporal data split, many users in the test set had no prior interactions in the training set.

Moreover, we observe different behavior across the three cities in terms of recommendation performance. In New York, the Pop recommender outperforms the rest, likely due to the limited data available, where popularity serves as a reliable proxy for user preferences. Consequently, this model exhibits the lowest scores in novelty and diversity. In contrast, in Tokyo, more sophisticated recommenders outperform the Pop approach for nDCG. Besides, the IB recommender achieves relatively low nDCG values but excels in novelty and diversity. Similarly, RP$^3\beta$ demonstrates strong performance in these dimensions; however, its results in nDCG are significantly lower. Regarding the other recommenders, we observe

Table 6: Performance of the recommenders in POI recommendation in terms of ranking accuracy (nDCG), novelty (EPC), and diversity (Gini) at cutoff 5. Best result for each metric in bold (excluding the Skyline).

City	NYC				TOK			
Method	nDCG	EPC	Gini	UC	nDCG	EPC	Gini	UC
Rnd	0.0000	**0.9981**	**0.4772**	**397**	0.0001	**0.9998**	**0.5710**	**26523**
Pop	**0.1096**	0.9385	0.0028	**397**	0.2260	0.8418	0.0001	**26523**
UB	0.0133	0.9927	0.1173	75	0.2343	0.8854	0.0024	19059
IB	0.0130	0.9945	0.1474	85	0.1422	0.9325	0.0674	19716
EASEr	0.0105	0.9890	0.0689	95	0.2081	0.8969	0.0014	19759
$RP^3\beta$	0.0106	0.9973	0.1639	95	0.0596	0.9794	0.1150	19759
BPR	0.0485	0.9425	0.0040	95	0.2350	0.8462	0.0001	19759
GeoBPR	0.0530	0.9670	0.0121	95	0.2323	0.8517	0.0001	19759
IRenMF	0.0327	0.9884	0.0949	95	**0.2390**	0.8487	0.0001	19759
H-PUM	0.1069	0.9437	0.0137	**397**	0.2190	0.8687	0.0170	**26523**
Skyline	0.8949	0.9826	0.0985	350	0.7707	0.9597	0.0623	26327

Table 7: Performance for route recommendation in terms of ranking accuracy (nDCG), novelty (EPC), diversity (Gini), and distance (Dist) of the route, at cutoff 5. Best result in bold.

City	Recommender	nDCG	EPC	Gini	Dist (km)	UC
NYC	Baseline-R	0.3584	0.9737	0.0078	0	16
	ClosestNN-R	0.4162	**0.9913**	**0.0385**	**0.544**	16
	MC-R	0.4285	0.9746	0.0138	0.905	16
	FMC-R	0.4130	0.9682	0.0070	5.637	16
	kNN-R	0.4253	0.9845	0.0131	1.226	16
	WG-R	**0.4332**	0.9880	0.0345	1.012	16
TOK	Baseline-R	0.3696	0.8555	0.0109	0	5870
	ClosestNN-R	0.3729	**0.9669**	**0.0392**	**0.053**	5870
	MC-R	**0.4250**	0.6954	0.0066	5.757	5870
	FMC-R	0.4110	0.7661	0.0024	15.939	5870
	kNN-R	0.4158	0.8683	0.0104	5.725	5870
	WG-R	0.4210	0.8076	0.0273	6.016	5870

that GeoBPR outperforms BPR in New York, whereas in Tokyo their performances are nearly equivalent. In terms of novelty and diversity, GeoBPR also achieves slightly better results, highlighting the critical role of geographical influence in POI recommendations. This observation is further supported by the performance of H-PUM and IRenMF models, with the latter achieving the highest ranking accuracy in 2 out of the 3 cities (regarding full results).

4.2.2 Route recommendation. Table 7 presents the performance for the route recommendation task. We observe that the MC-R model achieves the highest nDCG in Tokyo, and it is the second best in NYC, indicating superior ranking accuracy. Another well-performing recommender is the WG-R, being the best in NYC and the second best in the other city. In terms of novelty, diversity, and distance (EPC, Gini, Dist), the ClosestNN-R model performs best. This is because it relies solely on the distance component, ignoring both the characteristics of the POIs and the users' preferences. As a result, whenever there are multiple POIs located close to each other, the model prioritizes proximity, resulting in low Dist and high EPC and Gini values, but low accuracy. The results obtained by FMC-R are interesting, since it evidences lack of diverse recommendations. This may stem from its strategy of maximizing transitions between POI features, which can lead to repetitive cycles—transitions between similar POIs without exploring more varied options.

Table 8: Performance of the recommenders in contextual recommendation in terms of nDCG@5, when considering time, weather, or both as contexts. Best result for each context underlined, and the overall for each city in bold.

City	Recommender	Time	Weather	Both
NYC	C-Rnd	0.0018	0.0005	0.0005
	C-Pop	**0.0375**	<u>0.0060</u>	<u>0.0050</u>
	C-H-PUM	0.0254	0.0048	0.0045
TOK	C-Rnd	0.0000	0.0000	0.0000
	C-Pop	<u>0.0057</u>	**0.0113**	<u>0.0053</u>
	C-H-PUM	0.0048	0.0033	0.0031

4.2.3 Contextual recommendation. Table 8 shows the results after applying contextual post-filtering as described in Section 4.1.3. C-Pop is the best performing method independently of the context being used. However, it should be noted that in PTJ, C-Pop performs very well but is outperformed by C-H-PUM (see online appendix).

In NYC, the temporal context performs better, whereas in TOK the weather context produces better results. This is linked to Figures 4-6, where the weather seems more discriminative for the latter city. Finally, when compared against the results with no context (Table 6), it becomes evident that producing accurate contextual recommendations is much more difficult, since sparsity increases.

5 Conclusions and Future Work

Novel and flexible datasets are critical to advance research. In this paper, we have presented the *Context Trails* dataset, a new resource for the community to perform POI, route, and contextual recommendations by using geographical, categorical, weather, and temporal information. By extending a previous data source with newly collected information, we build a unique dataset with several contextual dimensions and the possibility to be used for POI or route recommendation.

Limitations. Routes included in *Context Trails* are limited to those in the original data [23], where coordinates were not available, thus requiring the use of Foursquare API (although the POI may not exist anymore). One possibility would be to start from the Foursquare Global Scale dataset [41] and generate routes from check-in data as in [30]. Moreover, LBSNs include very different types of POIs, not all of them might be interesting from a recommendation perspective (e.g., a bus stop); currently no filtering was done but included categories in the resource allow doing this systematically.

Future work. We aim to further expand our dataset with more cities (our current setup allows for this in a straightforward way) and more contexts (e.g., identifying events in those cities, like sport matches or concerts). It might also be useful to link this dataset with related datasets, such as Yelp or Gowalla, or more touristic oriented ones like tourist card logs [22]. Additionally, to allow finer analyses based on user types, we would like to discriminate users between tourists and locals, as done in recent research [11, 32].

Acknowledgments

Work supported by grant PID2022-139131NB-I00 funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe."

References

[1] Gediminas Adomavicius, Bamshad Mobasher, Francesco Ricci, and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. *AI Mag.* 32, 3 (2011), 67–80. https://doi.org/10.1609/aimag.v32i3.2364

[2] Jie Bao, Yu Zheng, David Wilkie, and Mohamed F. Mokbel. 2015. Recommendations in location-based social networks: a survey. *GeoInformatica* 19, 3 (2015), 525–565. https://doi.org/10.1007/s10707-014-0220-8

[3] Konstantin Bauman, Alexander Tuzhilin, and Moshe Unger. 2025. HyperCARS: Using Hyperbolic Embeddings for Generating Hierarchical Contextual Situations in Context-Aware Recommender Systems. *Inf. Syst. Res.* (2025), 1–25. https://doi.org/10.1287/isre.2022.0202

[4] Veronika Bogina, Tsvi Kuflik, Dietmar Jannach, Mária Bieliková, Michal Kompan, and Christoph Trattner. 2023. Considering temporal aspects in recommender systems: a survey. *User Model. User Adapt. Interact.* 33, 1 (2023), 81–119. https://doi.org/10.1007/s11257-022-09335-w

[5] Mario Casillo, Francesco Colace, Dajana Conte, Marco Lombardi, Domenico Santaniello, and Carmine Valentino. 2023. Context-aware recommender systems and cultural heritage: a survey. *J. Ambient Intell. Humaniz. Comput.* 14, 4 (2023), 3109–3127. https://doi.org/10.1007/s12652-021-03438-9

[6] Dawei Chen, Cheng Soon Ong, and Lexing Xie. 2016. Learning Points and Routes to Recommend Trajectories. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, Snehasis Mukhopadhyay, ChengXiang Zhai, Elisa Bertino, Fabio Crestani, Javed Mostafa, Jie Tang, Luo Si, Xiaofang Zhou, Yi Chang, Yunyao Li, and Parikshit Sondhi (Eds.). ACM, 2227–2232. https://doi.org/10.1145/2983323.2983672

[7] Lei Chen, Jie Cao, Haicheng Tao, and Jia Wu. 2023. Trip Reinforcement Recommendation with Graph-based Representation Learning. *ACM Trans. Knowl. Discov. Data* 17, 4 (2023), 57:1–57:20. https://doi.org/10.1145/3564609

[8] Lei Chen, Guixiang Zhu, Weichao Liang, and Youquan Wang. 2023. Multi-objective reinforcement learning approach for trip recommendation. *Expert Syst. Appl.* 226 (2023), 120145. https://doi.org/10.1016/J.ESWA.2023.120145

[9] Chen Cheng, Haiqin Yang, Irwin King, and Michael R. Lyu. 2012. Fused Matrix Factorization with Geographical and Social Influence in Location-Based Social Networks. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*, Jörg Hoffmann and Bart Selman (Eds.). AAAI Press. http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4748

[10] Visual Crossing Corporation. 2017–2018. Visual Crossing Weather. https://www.visualcrossing.com/. Accessed: November 2024.

[11] Linus W. Dietz, Rinita Roy, and Wolfgang Wörndl. 2019. Characterisation of Traveller Types Using Check-In Data from Location-Based Social Networks. In *Information and Communication Technologies in Tourism 2019, ENTER 2019, Proceedings of the International Conference in Nicosia, Cyprus, January 30-February 1, 2019*, Juho Pesonen and Julia Neidhardt (Eds.). Springer, 15–26. https://doi.org/10.1007/978-3-030-05940-8_2

[12] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati E. Pantziou. 2014. A survey on algorithmic approaches for solving tourist trip design problems. *J. Heuristics* 20, 3 (2014), 291–328. https://doi.org/10.1007/s10732-014-9242-5

[13] Sajal Halder, Kwan Hui Lim, Jeffrey Chan, and Xiuzhen Zhang. 2024. A survey on personalized itinerary recommendation: From optimisation to deep learning. *Appl. Soft Comput.* 152 (2024), 111200. https://doi.org/10.1016/J.ASOC.2023.111200

[14] Zheng Huang, Jing Ma, Yushun Dong, Natasha Zhang Foutz, and Jundong Li. 2022. Empowering Next POI Recommendation with Multi-Relational Modeling. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (Eds.). ACM, 2034–2038. https://doi.org/10.1145/3477495.3531801

[15] Sergio Ilarri, Raquel Trillo Lado, and Ramón Hermoso. 2018. Datasets for Context-Aware Recommender Systems: Current Context and Possible Directions. In *34th IEEE International Conference on Data Engineering Workshops, ICDE Workshops 2018, Paris, France, April 16-20, 2018*. IEEE Computer Society, US, 25–28. https://doi.org/10.1109/ICDEW.2018.00011

[16] Christopher Laß, Daniel Herzog, and Wolfgang Wörndl. 2017. Context-Aware Tourist Trip Recommendations. In *Proceedings of the 2nd Workshop on Recommenders in Tourism co-located with 11th ACM Conference on Recommender Systems (RecSys 2017), Como, Italy, August 27, 2017. (CEUR Workshop Proceedings, Vol. 1906)*, Julia Neidhardt, Daniel R. Fesenmaier, Tsvi Kuflik, and Wolfgang Wörndl (Eds.). CEUR-WS.org, 18–25. http://ceur-ws.org/Vol-1906/paper3.pdf

[17] Yang Li, Yadan Luo, Zheng Zhang, Shazia W. Sadiq, and Peng Cui. 2019. Context-Aware Attention-Based Data Augmentation for POI Recommendation. In *35th IEEE International Conference on Data Engineering Workshops, ICDE Workshops 2019, Macao, China, April 8-12, 2019*. IEEE, 177–184. https://doi.org/10.1109/ICDEW.2019.00-14

[18] Kwan Hui Lim, Jeffrey Chan, Shanika Karunasekera, and Christopher Leckie. 2019. Tour recommendation and trip planning using location-based social media: a survey. *Knowl. Inf. Syst.* 60, 3 (2019), 1247–1275. https://doi.org/10.1007/S10115-018-1297-4

[19] Yong Liu, Wei Wei, Aixin Sun, and Chunyan Miao. 2014. Exploiting Geographical Neighborhood Characteristics for Location Recommendation. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014, Shanghai, China, November 3-7, 2014*, Jianzhong Li, Xiaoyang Sean Wang, Minos N. Garofalakis, Ian Soboroff, Torsten Suel, and Min Wang (Eds.). ACM, 739–748. https://doi.org/10.1145/2661829.2662002

[20] Pablo Mateos and Alejandro Bellogín. 2025. A systematic literature review of recent advances on context-aware recommender systems. *Artif. Intell. Rev.* 58, 1 (2025), 20. https://doi.org/10.1007/S10462-024-10939-4

[21] Aditya Krishna Menon, Dawei Chen, Lexing Xie, and Cheng Soon Ong. 2017. Revisiting revisits in trajectory recommendation. In *Proceedings of International Workshop on Citizens for Recommender Systems, CitRec@RecSys 2017, 31 August 2017, Como, Italy*, Jie Yang, Zhu Sun, Alessandro Bozzon, Jie Zhang, and Martha Larson (Eds.). ACM, 2:1–2:6. https://doi.org/10.1145/3127325.3127326

[22] Sara Migliorini, Damiano Carra, and Alberto Belussi. 2021. Distributing Tourists among POIs with an Adaptive Trip Recommendation System. *IEEE Trans. Emerg. Top. Comput.* 9, 4 (2021), 1765–1779. https://doi.org/10.1109/TETC.2019.2920484

[23] Diego Monti, Enrico Palumbo, Giuseppe Rizzo, Raphaël Troncy, and Maurizio Morisio. 2018. Semantic Trails of City Explorations: How Do We Live a City. *CoRR* abs/1812.04367 (2018). arXiv:1812.04367 http://arxiv.org/abs/1812.04367

[24] Athanasios N. Nikolakopoulos, Xia Ning, Christian Desrosiers, and George Karypis. 2022. Trust Your Neighbors: A Comprehensive Survey of Neighborhood-Based Methods for Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, 39–89. https://doi.org/10.1007/978-1-0716-2197-4_2

[25] Bibek Paudel, Fabian Christoffel, Chris Newell, and Abraham Bernstein. 2017. Updatable, Accurate, Diverse, and Scalable Recommendations for Interactive Applications. *ACM Trans. Interact. Intell. Syst.* 7, 1 (2017), 1:1–1:34. https://doi.org/10.1145/2955101

[26] Hossein A. Rahmani, Mohammad Aliannejadi, Mitra Baratchi, and Fabio Crestani. 2020. Joint Geographical and Temporal Modeling Based on Matrix Factorization for Point-of-Interest Recommendation. In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12035)*, Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins (Eds.). Springer, 205–219. https://doi.org/10.1007/978-3-030-45439-5_14

[27] Hossein A. Rahmani, Mohammad Aliannejadi, Mitra Baratchi, and Fabio Crestani. 2022. A Systematic Analysis on the Impact of Contextual Information on Point-of-Interest Recommendation. *ACM Trans. Inf. Syst.* 40, 4 (2022), 88:1–88:35. https://doi.org/10.1145/3508478

[28] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*, Jeff A. Bilmes and Andrew Y. Ng (Eds.). AUAI Press, 452–461.

[29] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender Systems: Introduction and Challenges. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer, 1–34. https://doi.org/10.1007/978-1-4899-7637-6_1

[30] Pablo Sánchez and Alejandro Bellogín. 2020. Applying reranking strategies to route recommendation using sequence-aware evaluation. *User Model. User Adapt. Interact.* 30, 4 (2020), 659–725. https://doi.org/10.1007/s11257-020-09258-4

[31] Pablo Sánchez and Alejandro Bellogín. 2022. Point-of-Interest Recommender Systems Based on Location-Based Social Networks: A Survey from an Experimental Perspective. *ACM Comput. Surv.* 54, 11s (2022), 223:1–223:37. https://doi.org/10.1145/3510409

[32] Pablo Sánchez and Linus W. Dietz. 2022. Travelers vs. Locals: The Effect of Cluster Analysis in Point-of-Interest Recommendation. In *UMAP '22: 30th ACM Conference on User Modeling, Adaptation and Personalization, Barcelona, Spain, July 4 - 7, 2022*, Alejandro Bellogín, Ludovico Boratto, Olga C. Santos, Liliana Ardissono, and Bart P. Knijnenburg (Eds.). ACM, 132–142. https://doi.org/10.1145/3503252.3531320

[33] Diego Sánchez-Moreno, Álvaro Lozano Murciego, Vivian F. López Batista, María Dolores Muñoz Vicente, and María N. Moreno García. 2021. Dynamic inference of user context through social tag embedding for music recommendation. In *Proceedings of the 15th ACM Conference on Recommender Systems, RecSys, Workshop on Context-Aware Recommender Systems, 2021, Amsterdam, Netherlands, 27th September-1st October , 2021*. ACM, US.

[34] Harald Steck. 2019. Embarrassingly Shallow Autoencoders for Sparse Data. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 3251–3257. https://doi.org/10.1145/3308558.3313710

[35] Areej Bin Suhaim and Jawad Berri. 2021. Context-Aware Recommender Systems for Social Networks: Review, Challenges and Opportunities. *IEEE Access* 9 (2021), 57440–57463. https://doi.org/10.1109/ACCESS.2021.3072165

[36] Christoph Trattner, Alexander Oberegger, Lukas Eberhard, Denis Parra, and Leandro Balby Marinho. 2016. Understanding the Impact of Weather for POI Recommendations. In *Proceedings of the Workshop on Recommenders in Tourism co-located with 10th ACM Conference on Recommender Systems (RecSys 2016), Boston, MA, USA, September 15, 2016 (CEUR Workshop Proceedings, Vol. 1685)*, Daniel R. Fesenmaier, Tsvi Kuflik, and Julia Neidhardt (Eds.). CEUR-WS.org, 16–23. http://ceur-ws.org/Vol-1685/paper3.pdf

[37] Norha M. Villegas and Hausi A. Müller. 2010. Managing Dynamic Context to Optimize Smart Interactions and Services. In *The Smart Internet - Current Research and Future Applications*, Mark H. Chignell, James R. Cordy, Joanna Ng, and Yelena Yesha (Eds.). Lecture Notes in Computer Science, Vol. 6400. Springer, US, 289–318. https://doi.org/10.1007/978-3-642-16599-3_18

[38] Lin Wan, Han Wang, Yuming Hong, Ran Li, Wei Chen, and Zhou Huang. 2022. iTourSPOT: a context-aware framework for next POI recommendation in location-based social networks. *Int. J. Digit. Earth* 15, 1 (2022), 1614–1636. https://doi.org/10.1080/17538947.2022.2122611

[39] Hao Wang, Manolis Terrovitis, and Nikos Mamoulis. 2013. Location recommendation in location-based social networks using user check-in data. In *21st SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2013, Orlando, FL, USA, November 5-8, 2013*, Craig A. Knoblock, Markus Schneider, Peer Kröger, John Krumm, and Peter Widmayer (Eds.). ACM, 364–373. https://doi.org/10.1145/2525314.2525357

[40] Qingyun Wu, Huazheng Wang, Quanquan Gu, and Hongning Wang. 2016. Contextual Bandits in a Collaborative Environment. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*. ACM, US, 529–538. https://doi.org/10.1145/2911451.2911528

[41] Dingqi Yang, Daqing Zhang, and Bingqing Qu. 2016. Participatory Cultural Mapping Based on Collective Behavior Data in Location-Based Social Networks. *ACM TIST* 7, 3 (2016), 30:1–30:23. https://doi.org/10.1145/2814575

[42] Lina Yao, Quan Z. Sheng, Yongrui Qin, Xianzhi Wang, Ali Shemshadi, and Qi He. 2015. Context-aware Point-of-Interest Recommendation Using Tensor Factorization with Social Regularization. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 1007–1010. https://doi.org/10.1145/2766462.2767794

[43] Junjie Yin, Yun Li, Zheng Liu, Jian Xu, Bin Xia, and Qianmu Li. 2019. ADPR: An Attention-based Deep Learning Point-of-Interest Recommendation Framework. In *International Joint Conference on Neural Networks, IJCNN 2019 Budapest, Hungary, July 14-19, 2019*. IEEE, 1–8. https://doi.org/10.1109/IJCNN.2019.8852309

[44] Fajie Yuan, Joemon M. Jose, Guibing Guo, Long Chen, Haitao Yu, and Rami Suleiman Alkhawaldeh. 2016. Joint Geo-Spatial Preference and Pairwise Ranking for Point-of-Interest Recommendation. In *28th IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2016, San Jose, CA, USA, November 6-8, 2016*. IEEE Computer Society, 46–53. https://doi.org/10.1109/ICTAI.2016.0018

[45] Jiyong Zhang, Xin Liu, Xiaofei Zhou, and Xiaowen Chu. 2021. Leveraging graph neural networks for point-of-interest recommendations. *Neurocomputing* 462 (2021), 1–13. https://doi.org/10.1016/J.NEUCOM.2021.07.063

[46] Ruining Zhang, Chanjuan Liu, Qiang Zhang, and Xiaopeng Wei. 2023. A Heuristic Framework for Personalized Route Recommendation Based on Convolutional Neural Networks. In *PRICAI 2023: Trends in Artificial Intelligence - 20th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2023, Jakarta, Indonesia, November 15-19, 2023, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 14327)*, Fenrong Liu, Arun Anand Sadanandan, Duc Nghia Pham, Petrus Mursanto, and Dickson Lukose (Eds.). Springer, 269–280. https://doi.org/10.1007/978-981-99-7025-4_24

[47] Shiming Zhang, Zhipeng Luo, Li Yang, Fei Teng, and Tianrui Li. 2024. A survey of route recommendations: Methods, applications, and opportunities. *Inf. Fusion* 108 (2024), 102413. https://doi.org/10.1016/J.INFFUS.2024.102413

[48] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Jiajie Xu, Zhixu Li, Fuzhen Zhuang, Victor S. Sheng, and Xiaofang Zhou. 2022. Where to Go Next: A Spatio-Temporal Gated Network for Next POI Recommendation. *IEEE Trans. Knowl. Data Eng.* 34, 5 (2022), 2512–2524. https://doi.org/10.1109/TKDE.2020.3007194

[49] Benyou Zou, Cuiping Li, Liwen Tan, and Hong Chen. 2015. GPUTENSOR: Efficient tensor factorization for context-aware recommendations. *Inf. Sci.* 299 (2015), 159–177. https://doi.org/10.1016/j.ins.2014.12.004