Proceedings of the Seventeenth International AAAI Conference on Web and Social Media (ICWSM 2023)

Predicting Future Location Categories of Users in a Large Social Platform

Raiyan Abdul Baten[1*], Yozen Liu[2], Heinrich Peters[3], Francesco Barbieri[2], Neil Shah[2], Leonardo Neves[2], Maarten W. Bos[2]

[1]Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14623, USA
[2]Snap Inc., Santa Monica, CA 90405, USA
[3]Columbia University in the City of New York, New York, NY 10027, USA
rbaten@ur.rochester.edu, {yliu2, fbarbieri, nshah, lneves, maarten}@snap.com, heinrichpeterz@gmail.com

Abstract

Understanding the users' patterns of visiting various location categories can help online platforms improve content personalization and user experiences. Current literature on predicting future location categories of users typically employs features that can be traced back to the users, such as spatial coordinate histories and demographic identities. Moreover, existing approaches commonly suffer from cold-start and generalization problems, and often cannot specify when the user will visit the predicted location category. In a large social platform, it is desirable for prediction models to avoid using user-identifiable data, generalize to unseen and new users, and be able to make predictions for specific times in the future. In this work, we construct a neural model, LocHabits, using data from Snapchat. The model omits user-identifiable inputs, leverages temporal and sequential regularities in the location category histories of Snapchat users, and predicts the users' next-hour location categories. We evaluate our model on several real-life, large-scale datasets from Snapchat and FourSquare, and find that the model can outperform baselines by 14.94% accuracy. We confirm that the model can (1) generalize to unseen users from different areas and times, and (2) fall back on collective trends in the cold-start scenario. We also study the relative contributions of various factors in making the predictions and find that the users' visitation preferences and most-recent visitation sequences play more important roles than time contexts, same-hour sequences, and social influence features.

Introduction

A key characteristic of online social platforms and content distributors is their ability to adapt to the user. Unlike traditional mass media which cater to a broad target audience, online platforms can tailor their contents, interactive features, and ads in a way that is personalized, valuable, and appropriate for the user (Berkovsky, Kuflik, and Ricci 2008). A prerequisite for successful personalization is to predict user activities ahead of time: knowing when a user will go shopping, get coffee, exercise, attend a lecture, or visit an amusement park, for instance, can allow the platform to deliver the right content and user experience at the right time. Importantly, people's activities can be strongly tied to the locations they visit, e.g., going to a gym for exercise or to a French restaurant to eat French food (Yang et al. 2014). Thus, predicting the *location categories* people will visit in the future (henceforth referred to as 'future location categories') can provide a strong marker of people's activities at that time, and in turn, provide contextual information for rich personalization.

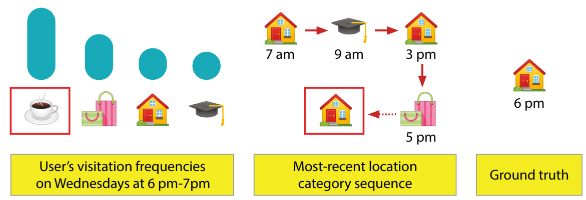

Figure 1: Illustration of the task. We wish to predict a user's location category on Wednesday between 6 pm-7 pm, using data up to 6 pm. Based on the user's historical visitation frequencies at the query hour, 'Beverages' can be guessed as the most likely location category (red box in the *left* panel). Again, based on the most recent location category sequence, 'Residence' can be predicted instead (*center*). Given the ground truth location category (*right*), an end-to-end model can be trained that makes the next-hour prediction by fusing information from such modalities.

Humans show surprising regularity in their individual and collective mobility patterns despite enjoying considerable freedom in making their movement and destination choices (Song et al. 2010). Such regularity can partly be explained by people's 'stable habits', where they carry out routine activities under routine situational cues such as the same time, place, or context (Aarts and Dijksterhuis 2000). As a result, regularities in people's past visitation histories can be used to predict their future location categories. A growing body of work has focused on predicting people's future location categories directly (Liao et al. 2018a; Liu et al. 2019). Additionally, a tangential direction has sought to predict people's precise location coordinates or Points of Interest (POIs) (Yang et al. 2019), from which the categories can be inferred. The majority of these works use timestamped location coordinates to mine spatiotemporal regularities. Demographic features such as gender/age (Chang and Sun 2011) and visitations of social ties and non-social

[*]Work done when the author was an intern at Snap Inc.

co-locators (Chen et al. 2022) have also been shown to hold information about people's future whereabouts.

Existing challenges and our solutions. Several considerations arise in implementing a prediction framework for the said purpose in a real-world, global-scale online social platform that have not been adequately addressed in literature:

- **Avoiding privacy-intrusive and user-identifiable data:** Demographic features such as age, race, and gender are not only privacy-intrusive but also static and exclusionary. One of the most informative features, spatial coordinates, is also privacy sensitive, as the data can potentially be traced back to a user (De Montjoye et al. 2013).
  As such, we avoid using data that may identify the users in this work. We instead leverage temporal regularities in the location category histories of users and their friends to make future location category predictions.
- **Making predictions at a specific future time:** Most of the current literature focuses on predicting the next location or the next location category of a user, ignoring when it will occur (Chang and Sun 2011). This limitation makes the predictions less actionable, as the visitations can happen at any arbitrary time. Some recent works have acknowledged and addressed this issue: e.g., Yu et al. predicted a user's most likely locations in the next 24 hours (Yu et al. 2020), and Karushima et al. predicted future user actions along with the timing (Kurashima, Althoff, and Leskovec 2018). The trade-off consideration here is, if the predictions are made too late, the platform will not have enough time to act on the insight; but if it is made too long in advance, the prediction will likely miss out on information from the more recent user activities and thus suffer from inaccuracy.
  In this work, we focus on making location category predictions for the next hour. This specification enables the platform to proactively personalize contents and features ahead of time while allowing the model to leverage the users' activities up to the latest hour in making predictions. We unearth novel insights on how best to extract temporal patterns from location category histories to make next-hour predictions.
- **Addressing the generalization and cold start problems:** Many previous approaches require learning user-specific preferences (Yang et al. 2014; Wu et al. 2022). In platforms with hundreds of millions of users, however, learning parameters for each user does not scale well. Again, models that employ one-hot encoded representations of users (e.g., (Yao et al. 2017)) cannot be readily applied to a new/unseen user. Instead, a single trained model is desirable which generalizes to out-of-sample users, including users of different areas (cities/states/countries) and times. In cold start situations (Schein et al. 2002) with little to no data from a new user, the model should be able to fall back on collective trends for making predictions.

Given these considerations, we build a deep neural model, LocHabits, using timestamped location category and social graph data from Snapchat. The model takes a modular approach and fuses information from sub-networks capturing the users' visitation frequencies, their friends' aggregated visitation frequencies, time contexts, same-hour location category sequences, and most-recent location category sequences to make the next-hour category predictions (Figure 1). We conduct extensive experiments to evaluate our model's performance in several large-scale real-life dataset versions from Snapchat and FourSquare. We find that the model outperforms comparison baselines by 14.94% on average in terms of prediction accuracy, while gracefully generalizing to new and unseen users from different areas and times. We study the relative contributions of the different modules and find the users' visitation frequencies and the most recent location category sequences to be the most predictive features.

Summary. Our contributions can be summarized as:

- We take a privacy-first approach to predict people's next-hour location categories in a large-scale industrial setting.
- We unearth novel insights on how best to extract temporal patterns from people's past location category histories.
- We build the LocHabits prediction model and confirm the generalization of the model to new and unseen users in datasets from various geographic areas and time periods.
- We shed light on the relative contributions of various prediction features, and find the users' visitation frequencies and the most recent sequences to be the most predictive.

Related Work

Predicting User Activities

The user activity prediction literature typically tackles two kinds of problems: predicting *in-app* and *offline* activities. For instance, it is known that people's ephemeral (most recent) and cyclical (e.g., weekday vs weekend habits) in-app activities can help predict their future in-app activities (Chowdhury et al. 2021).

Our focus in this paper is to capture offline activities through location categories. To this end, various models have been proposed that predict people's current or future location categories. These models typically employ user IDs, spatial coordinates, timestamps, and location category data as inputs. For example, Cui et al. predicted the current activity of the users as they tweeted, using an LSTM architecture (Cui, Agrawal, and Ramnath 2020). To et al. used social, spatial, temporal, and semantic data to predict a traveler's next activity using a neural model (To, Si, and Chen 2019). Liao et al. captured the interplay between spatial locations and activities using an RNN-based model, underpinned by the sequential dependencies and temporal regularities of spatial-activity topics (Liao et al. 2018a). For recommending activities to users, approaches like a collaborative tensor-topic factorization model (Liu et al. 2019) and a personalized time-aware collaborative model (Rahimiaghdam, Karagoz, and Mutlu 2016) have been proposed. Zhang et al. adopted an NLP-based approach (i.e., using n-gram and PLSA) to capture long-range dependencies and contextual factors to predict future location categories (Zhang et al. 2017). Yang et al. (Yang et al. 2014) noted that the users' mobility is usually confined to specific geographic regions ('Personal Functional Regions'), which they used to model the spatial preferences of users. The temporal preferences were captured collaboratively, and the two modalities were fused together for predicting future location categories. Lian et al. clustered users using spatial, temporal, and sequence features, and collaboratively predicted future location categories (Lian and Xie 2011). Ye et al. built a mixed hidden Markov model that leveraged sequential patterns from location categories, as well as spatial and temporal contexts, to predict a user's next location category (Ye, Zhu, and Cheng 2013). Using a multivariate temporal point processes model, Kurashima et al. showed that user preferences, time-varying propensities of actions, short-term dependencies between actions, and long-term periodic effects hold information for predicting future actions of users as well as the associated timing (Kurashima, Althoff, and Leskovec 2018).

Predicting Points of Interest

A growing body of literature has tackled the problem of predicting Points of Interest (POI). Geo-location coordinates are among the defining components of a POI, which makes the prediction space sparser than predicting location categories. This body of work varies in terms of the (1) model specifications, (2) problem formulations, and the (3) influencing factors incorporated, as elaborated below:

(1) Model specifications. Earlier attempts used Matrix Factorization (Rahmani et al. 2020), Bayesian Personalized Ranking (He, Li, and Liao 2017), Collaborative filtering (Qiao et al. 2018), and Markov Chain (Chen, Liu, and Yu 2014) based models, among other classical machine learning approaches. Recently, a large variety of deep learning-based models have been proposed. Notably, Liu et al. extended RNN with spatiotemporal information in the Spatial-Temporal RNN (ST-RNN) model (Liu et al. 2016). The Attentional Recurrent Neural Network (ARNN) (Guo et al. 2020) captured both sequential and transitional regularities. In general, these deep models leverage CNN (Elmi, Benouaret, and Tan 2021), RNN (Liao et al. 2018b; Yang et al. 2020), LSTM (Yu et al. 2020; Yao et al. 2017; Cui et al. 2021), self-attention (Guo et al. 2020; Wang et al. 2021), graph embedding (Xie et al. 2016; Christoforidis et al. 2018), and relevant technologies (Islam et al. 2022).

(2) Problem formulations. There are variations in the tasks these models aim to perform. Some approaches recommend *unvisited* POIs to people (Christoforidis et al. 2021), while others predict people's most likely POIs based on their previous visitations (Doan, Yang, and Reddy 2019). Given a user's sequence of most recent POIs, many models predict only the next POI in the sequence (sequence to one) (Wu et al. 2022) or successive upcoming POIs (sequence to sequence) (Yu et al. 2020; Cheng et al. 2013; Chang et al. 2018). For instance, a variational attention-based model was used to capture the sequence information from spatial and temporal data for predicting the next POIs (Gao et al. 2019). While most models predict next/successive POIs ignoring the time of those event(s), some models explicitly consider the query time. For example, the Category-Aware Deep Model (CatDM) (Yu et al. 2020) attempts to predict POIs likely to be visited in the next 24 hours. In doing so, the authors divide each day into 12 time periods and the days in the week into weekdays and weekends. Cao *et al.* used spatial, temporal, and graph data to predict the users' locations at any fine-grained future hour (Cao et al. 2018). Burbey applied a Markov Model to timestamped spatial data to predict location at a specific future time (Burbey 2011). Yang *et al.* proposed LBSN2Vec, which takes a hypergraph embedding approach to predict future POIs (Yang et al. 2019).

(3) Influencing factors. The proposed models also differ in the data they use and the features they extract from the data. As spatial proximity can dictate where people will go next, most models leverage *geographical influence* in making the predictions (Sun et al. 2020; Chang et al. 2018; Kefalas and Manolopoulos 2017; Li et al. 2021). A large body of works has shown that *sequential effects* play a large role in predicting POIs: much like a sentence, where the sequence of previous words can hold information about the upcoming word (Guo et al. 2020; Wu et al. 2022; Kurashima, Althoff, and Leskovec 2018). *Temporal influence* has also been shown to be a prominent predictor (Cao et al. 2018; Doan, Yang, and Reddy 2019; Li, Shen, and Zhu 2018; Li et al. 2021; Kurashima, Althoff, and Leskovec 2018) since human activities can be strongly associated with the hour of the day and the day of the week (Chang and Sun 2011). Furthermore, the locations of one's social ties and non-social co-locators have proven to contain important information regarding people's whereabouts (Chen et al. 2022; Chang and Sun 2011). For instance, long-distance travel is often influenced by a friend living there (Cho, Myers, and Leskovec 2011). Thus, *social influence* features can be incorporated to improve POI predictions. User attributes such as *demographic features* (age/gender) have additionally been used to predict where people will go next (Chang and Sun 2011). The *location categories* can establish the contexts for making POI predictions (Yu et al. 2020; Wu et al. 2022). Furthermore, the *historical preferences* of the users (i.e., previous visitation counts) are known to be significant predictors of future POIs (Gao et al. 2019; Kurashima, Althoff, and Leskovec 2018). Recently, it has been shown that people's destination choices can be impacted by the choice of the location-based service itself (Ochiai et al. 2020). Several other works leveraged features of weather (Nawshin et al. 2020) and textual sentiment of reviews (Chang et al. 2018) to make the predictions.

Notably, the bulk of user activity and POI prediction literature utilizes spatial geo-location coordinate data. In contrast, we drop spatial coordinates and user-identifiable sources and use location categories, timestamps, and social graph data (detailed in the next section). This makes our problem scope different than most published works. Using these data, we curate feature sets as informed by previous works as well as our own explorations detailed in the sequel.

Constructing the Datasets

Privacy Considerations

For illustrative purposes, contrast the two examples below.

Example 1. A user (female, aged 23) visited the auditorium of the University of X (latitude $\underline{a_1}$, longitude $\underline{b_1}$) at 8 am on Monday. She then visited Joe's Coffee Shop ($\underline{a_2}$, $\underline{b_2}$) at 1 pm. Later, she went home ($\underline{a_3}$, $\underline{b_3}$) at 5 pm.

Figure 2: A tag cloud of the location categories.

Dataset	Time Span	Location	# Users
Sn-CA-Jun	Jun 1-Jun 30, 2021	California	210k
Sn-CA-Aug	Jul 29-Aug 28, 2021	California	224k
Sn-NY-Aug	Jul 29-Aug 28, 2021	New York	93k
Sn-TX-Aug	Jul 29-Aug 28, 2021	Texas	209k
Fs-US-Apr	Apr 3-Jul 11, 2012	USA	13278
Fs-JP-Apr	Apr 3-Jul 12, 2012	Japan	6802
Fs-BR-Apr	Apr 3-Jul 12, 2012	Britain	8878

Table 1: Dataset Summary. We use four dataset versions from Snapchat (Sn) and three from FourSquare (Fs).

Example 2. A user visited the location category "University" at 8 am on Monday, followed by "Beverages" at 1 pm, and finally the "Residence" category at 5 pm.

Given sufficient data, a malicious actor might be able to trace the user in Example 1 using the gender, age, and spatiotemporal information (underlined). In fact, only four spatiotemporal points are enough to uniquely identify 95% of individuals (De Montjoye et al. 2013), although methods like differential privacy can offer increased protection (Pyrgelis, Troncoso, and De Cristofaro 2017). In Example 2, only high-level categories are used. This sequence can belong to a large set of candidate users from around the globe, making it harder to pinpoint the user. Intuitively, using non-personally identifiable and high-level data can offer superior protection of one's privacy.

Building on these intuitions, we limit our scope to predicting future location categories only from (1) previous location categories, (2) timestamps, and (3) aggregated information on location categories from one's friends. Using friends' information merits caution: for example, if I have only one friend, and that friend's visitation history is used to predict my future location category and in turn personalize the content delivered to me, it can compromise the friend's privacy. It is therefore important to ensure that the users have at least a large threshold number of friends, that the information from those friends is aggregated appropriately, and that all of the involved users provided consent. We avoid using (1) spatial coordinates, (2) demographic information (age, gender, race, etc.), (3) weather (which can help narrow down candidate locations), (4) textual reviews (which can help to look up specific venues), etc. data that are frequently used in the literature. We acknowledge that the accepted and omitted lists of data sources are not comprehensive, and other sources that are non-personally identifiable and are sufficiently high-level can be used as well.

Dataset Curation

We conduct our explorations primarily on Snapchat's location visitation data. Snapchat is an online social and instant messaging platform. With the users' permission[1], Snapchat senses precise location coordinates using methods that include GPS, wireless networks, cell towers, Wi-Fi access points, and other sensors when the app is opened. Using internal methods, the coordinates are mapped to the most probable venues or business names (e.g., Joe's Coffee Shop), and, in turn, to the respective low-level location category names (e.g., Coffee shop). We use the venue categories from Foursquare[2]—where the categories are organized in hierarchical chains—to convert the low-level category names to high-level categories (as available at the time of analysis). We use the second-highest level from these chains for our analysis, or the first level if there is only one category in the hierarchy. For example, in the hierarchical chain "Arts & Entertainment > Stadium > Baseball Stadium", we take "Stadium" as the location category of interest (see Figure 2).

We collect 4 anonymized dataset versions from Snapchat (Table 1). All of these versions are from the United States. The first version is from users in California, USA, spanning the month of June 2021. We collect three more versions in August 2021, from users in California, New York, and Texas—three states geographically located at different ends (west, east, and south) of the USA. This allows us to test our model's performance on users of the same and different states, at different times.

We only have glimpses of the users' true visitation histories, based on whether they opened Snapchat during their visits. To mitigate potential errors from unobserved data, we analyze only the most active users, as recommended in previous literature (Song et al. 2010). To that end, we select users who logged at least 30 location entries in the month. We also collect friendship graphs from each dataset version, where each user has at least a threshold number of friends.

To further validate the generalizability of our findings, we use FourSquare check-in data from the USA, Japan, and Britain to get three geographically separated dataset versions, as constructed from the publicly available global-scale dataset in (Yang et al. 2019) (Table 1). The FourSquare dataset versions come with pre-processed low-level location category names, which we convert to the higher-level location categories using the same procedure as the Snapchat datasets. The friendship graphs among the users in these dataset versions are crawled from Twitter by the authors. We restrict our analysis to the data from the first 100 days in these dataset versions since the per-day data volume becomes relatively sparse afterward.

People's location visitation data is known to be heavy-

[1]https://snap.com/en-US/privacy/privacy-policy

[2]https://developer.foursquare.com/docs/build-with-foursquare/categories/

tailed (Gao, Tang, and Liu 2012), which we confirm in all of the dataset versions. From 141 location categories in the Snapchat datasets, we limit our analyses to the 50 most frequent categories, to avoid spending computational resources on predicting categories that the users rarely ever visit and/or log. These 50 categories comprise nearly all of the Snapchat data (99.29%). We limit our analyses to the 50 most frequent location categories in the FourSquare datasets as well, which comprise 98.53% of all its data.

Task Description

Predicting Next-hour Location Categories

Some applications may require predicting a user's locations/categories in the near future (e.g., route planning for taxi drivers to maximize potential passenger pickups (Yang et al. 2017)). Again, some applications (e.g, traffic/infrastructure planning) may need to predict locations long in advance by leveraging long-term periodicity (Kurashima, Althoff, and Leskovec 2018). Unfortunately, most frameworks predict the next location without specifying when it will occur. For our use-case of content personalization in online social platforms, at any given hour in the day, we wish to predict the location categories of the users in the upcoming hour. This choice leaves enough time margin for the platform to curate the right user experiences while making use of the users' recent visitation information from up to the previous hour.

Problem definition. More formally, let us consider a set of users U. Each user $u \in U$ visits locations described by (i) location categories $c \in C$, and (ii) timestamps in terms of the day of the week d and hour in the day h (discrete 1-hour bins). Thus, we have a set of user visitation sequences $S = \{s_1, s_2, ..., s_{|U|}\}$, where each sequence $s_u = v_1, v_2, ..., v_{m_u}$ consists of m_u chronologically ordered visitation records from user u. Each record is a tuple in the form of $v = (c, d, h)$. The undirected friendship ties $e \in E$ among the users are captured in graph $\mathcal{G}(U, E)$. Let, $G(u) = \{v : e_{uv} \in E\}$ denote the set of direct friends of user u. Our task can then be expressed as,

PROBLEM 1: PREDICTING THE USER'S LOCATION CATEGORY IN THE NEXT HOUR. *Learn a model from S and $\mathcal{G}$, that predicts the probability of each location category $c \in C$ to be visited by $u \in U$ at the query day and hour (d, h), using the s_u and $s_{v \in G(u)}$ data up to the previous hour of the query.* In other words, predict

$$Pr(c_u | d, h, s_u, s_{v \in G(u)}). \quad (1)$$

The output is a ranked list of k location categories with the highest predicted probabilities.

Evaluation Strategy

Each of the Snapchat datasets has one month of data. We take the first three weeks as the observation period to build visitation histories of the users, and the fourth week as the prediction period. For example, if a user has 30 entries in the observation period and 8 entries in the prediction period, we construct 8 data points from the user with the prediction period entries as ground truth target labels. For the FourSquare datasets, we have 100 days of data, which we split into 70 days of observation period and 30 days of prediction period. The data points are randomly divided into 60:10:30 splits as the training, validation, and test sets while ensuring that data from the same user remains in the same split to avoid information leakage.

Since we discretize time into 1-hour bins, there can be multiple correct categories for a query at (d, h) if a user visits multiple categories within the hour. We take the set of correct categories as the 'relevant' items. We return a sorted list of k categories as the 'recommended' items. From the 'relevant' and 'recommended' lists, we compute ranking performance using popular metrics: (1) Normalized Discounted Cumulative Gain (NDCG@k), (2) Recall@k, (3) Reciprocal Rank (RR@k), (4) Coverage Error (CE), and (5) Macro-F1.

Exploratory Analysis

In this section, we explore the datasets' predictive capabilities using entropy-based and heuristic approaches. The objective is to gather insights on informative features that can later be used in the more robust neural model.

Entropy Analysis

Entropy is arguably the most fundamental quantity that captures predictability in time-series data (Navet and Chen 2008). To quantify the interplay between regular (thus predictable) and random (thus unforeseeable) trends at individual-level time-series data, here we contrast three entropy measures (Song et al. 2010):

- **Random entropy:** If user u visits n_u unique location categories, then a naive algorithm can randomly pick any of the n_u categories with equal probability as the user's predicted next-hour venue. The uncertainty in the prediction can then be quantified as $S_u^{\text{rand}} = \log_2 n_u$.
- **Shannon entropy:** If we know the frequency distribution of user u's visits to each of the n_u categories, we can leverage the heterogeneity of visitation patterns in making the prediction. The entropy is then given by $S_u^{\text{unc}} = -\sum_{z=1}^{n_u} p_u(z) \log_2 p_u(z)$, where $p_u(z)$ is the historical probability of user u visiting category z.
- **Lempel-Ziv data compression:** If we know the frequency distribution as well as the sequence of the visited categories of the user u, we can quantify the 'true' uncertainty as, $S_u^{\text{lz}} = (L_u \log_2 L_u) / \sum_{l=1}^{L} \Lambda_l$, where L_u is the total length of user u's historical time series, and Λ_l is the shortest subsequence in the time-series that starts at position l and doesn't previously appear from position 1 to $l-1$ (Song et al. 2010).

Heuristic Approaches for Prediction

- **Overall collective preferences** (Doan, Yang, and Reddy 2019; Ye, Zhu, and Cheng 2013): We sort the location categories based on the collective visitation frequencies (i.e., sort by collective popularity). This sorted list is returned in response to every query.
- **Collective preferences at hour h:** We return a sorted list of categories based on collective-level frequencies at the query hour h.

Dataset	Entropy			Collective preferences		User preferences			
	S^{rand}	S^{unc}	S^{lz}	Overall	At hour h	Overall	Day/night splits	Composite Q_c	O-z MM
`Sn-CA-Jun`	3.04±0.5	2.19±0.5	2.09±0.4	24.88%	27.60%	48.06%	48.79%	**49.26%**	38.15%
`Sn-CA-Aug`	3.02±0.5	2.17±0.5	2.07±0.4	22.45%	27.89%	46.94%	51.34%	**52.19%**	40.00%
`Sn-NY-Aug`	2.97±0.5	2.11±0.6	2.03±0.4	24.31%	25.02%	46.95%	48.93%	**49.38%**	44.17%
`Sn-TX-Aug`	3.02±0.5	2.15±0.5	2.04±0.4	29.36%	33.95%	47.13%	52.47%	**53.34%**	42.12%
`Fs-US-Apr`	3.53±0.5	2.77±0.5	2.11±0.6	25.09%	28.43%	32.35%	35.04%	**36.50%**	25.59%
`Fs-JP-Apr`	3.23±0.5	2.27±0.5	1.74±0.6	32.28%	38.06%	50.88%	51.61%	**53.07%**	44.00%
`Fs-BR-Apr`	3.49±0.5	2.74±0.5	1.86±0.6	18.58%	26.72%	31.67%	35.31%	**37.51%**	24.92%

Table 2: Results of the exploratory analyses. The entropy values (mean±SD) are given in bits. The collective preference and user preference-based predictions are given in NDCG@1 percentages. The dataset-wise best predictions are highlighted in bold.

- **Overall user preferences** (Noulas et al. 2012; Ye, Zhu, and Cheng 2013): We sort location categories based on a user's overall visitation counts. This list is returned in response to every query from that user.
- **User preferences under day-time/night-time splits:** We split the data into 'day-time' and 'night-time' segments, to capture trends strictly associated with those splits. Our extensive explorations show that if the night-time range is too narrow, it results in too few data points in the night split and thus leads to poor prediction performances. We find that taking 10 am to midnight as the 'day'-time and the rest of the hours as 'night'-time gives the best performances. We thus create two sorted lists for each user based on their visitation frequencies in those two time-splits. The appropriate list is returned for the query hour.
- **User preferences using a composite score:** Here we combine information from a user's overall, day-based, and hour-based visitation counts into one score. Namely, we first split the user's data into daytime and nighttime-based data frames as described above. From these splits, we collect the user's total ($count_{tot}$), day-based ($count_d$), and hour-based ($count_h$) visitation counts for each location category. We sum these counts to compute a composite score for each category c at a given time (d, h) as, $Q_c(d, h) = count_{tot} + count_d + count_h$. We sort the categories based on their composite scores and return the sorted list as the query response.
- **Order-z Markov Model** (Doan, Yang, and Reddy 2019; Yang et al. 2014; Ye, Zhu, and Cheng 2013): We use order-z Markov models with $z = \{1, ..., 5\}$ to generate sequence-based prediction baselines. For example, in an order-4 Markov model, we use a sequence of 4 previous location categories to predict the next category with the maximum likelihood. Our explorations indicate that the best results are obtained for $z = 3$, which we report here.

Analysis Results

We compute the three entropy measures from each user's time-series data and report the mean and SD of the collective distributions. Table 2 shows the results. In all of the datasets, we have $S^{\text{lz}} < S^{\text{unc}} < S^{\text{rand}}$, where the differences are statistically significant in Mann-Whitney U-test ($p < 10^{-4}$ in each case). This gives us multiple insights:

1. While there are 50 location categories in our datasets, at an individual level, a user only visits a small number of those categories. This is reflected in $S^{\text{rand}} \approx 3.19$ bits, suggesting that a user who randomly chooses their next category can be found on average in any of $2^{3.19} = 9.1$ location categories out of 50.
2. If one's heterogeneity in visitation patterns is incorporated (i.e., in S^{unc}), the uncertainty drops, and a typical user can be found in one of $2^{2.34} \approx 5.1$ categories. This suggests that user-level visitation *frequencies* contain substantial information for making predictions.
3. Finally, the further reduced S^{lz} values suggest that the user-level visitation *sequences* contain additional information in making the predictions, as, in this case, the user can be found in any of $2^{1.99} \approx 4$ location categories.

Table 2 also lists the prediction performances of the heuristic approaches in terms of NDCG@1 (i.e., the accuracy of the best guess being correct). Given the skewed distribution of the category frequencies, a naive baseline is given by choosing the majority category every time (i.e., by the $k = 1$ output from the overall collective preferences heuristic), which results in $\approx 25.3\%$ accuracy. Collective preferences at hour h lead to better predictions than time-agnostic overall collective preferences. The user-preference-based predictions comfortably outperform the collective-preference-based ones, which is expected based on the insights from the entropy analysis above. Furthermore, binary-splitting the data frames based on day-time and night-time consistently gives better results than using overall visitation frequencies. Using the composite scores gives the best results among these heuristic approaches. This suggests that combining information from overall, day-based, and hour-based splits of the data is promising for making superior predictions. The order-3 Markov model outperforms the overall collective preferences baseline, showing promise for the use of *sequential* modeling of location category data and corroborating the entropy analysis results. We use these insights to build the deep neural network in the next section.

The LocHabits Neural Model

Below, we describe LocHabits, a neural model for predicting the users' next-hour location categories.

Feature Vectors and Target Labels

Our target output is a $\mathbb{R}^D$ vector that captures the predicted probability of each location category, where $D = |C|$ is the number of categories. We take a modular approach in building the network, where 5 feature vectors are processed in 5 different modules or sub-networks before being fused together via an attention mechanism. We construct separate feature vectors for the modules $\{T, I, F, M, L\}$, where T= time context, I = user's visitation frequencies, F = friends' aggregated visitation frequencies, M = most-recent location category sequence, and L = same-hour location category sequence. This modular approach allows us to (1) process each feature branch suitably to extract the information therein (e.g., using LSTMs for sequences), (2) fuse the modules together to capture their interactions, and (3) gain insights into the relative contributions of the modules through attention weights. Below we describe the feature vectors in detail:

Time context. The day of the week and the hour in the day are known to be highly informative context features of people's whereabouts (Chang and Sun 2011). Capturing the joint dynamics of the time contexts and the associated location categories can help prioritize relevant categories based on the query day and hour. When trained over data from thousands of users, this module helps capture the collective preferences associated with the respective time contexts. It is common in the literature to discretize time in the day into 12 (Yu et al. 2020) or 24 bins (Li, Shen, and Zhu 2018; Wu et al. 2022); and the day in the week into 2 (weekday and weekends) or 7 bins (Li, Shen, and Zhu 2018; Yu et al. 2020). We take one-hot encoded vectors to represent the day of the week, $Z_T^d \in \mathbb{R}^{D_d}$, and hour of the day, $Z_T^h \in \mathbb{R}^{D_h}$, where $D_d = 7$ and $D_h = 24$. The two vectors are concatenated to give the time context feature vector, $Z_T = Z_T^d \oplus Z_T^h$, where $\oplus$ represents concatenation, and $Z_T \in \mathbb{R}^{D_d + D_h}$.

User's visitation frequencies. Our exploratory analysis showed promise in employing user-level visitation frequencies as a predictor. To capture user preferences, we compute three 'count' vectors based on the visitation data of each user. We first split the data frames into day-time and night-time segments, using the same time-split boundaries as informed by our exploratory analysis. $Z_I^{tot} \in \mathbb{R}^D$ holds the user's max-min normalized overall visitation counts at the D categories, from the day-time or the night-time split corresponding to the query hour. From the same day-time/night-time splits, we also collect $Z_I^d \in \mathbb{R}^D$ and $Z_I^h \in \mathbb{R}^D$, which respectively hold the max-min normalized visitation counts from the same day of the week d and hour in the day h as the query. We concatenate these three count vectors to get a user-level visitation frequency feature vector, $Z_I = Z_I^{dn} \oplus Z_I^d \oplus Z_I^h$, where $Z_I \in \mathbb{R}^{3D}$.

Friends' aggregated visitation frequencies. It is known in the literature that the friends' visitation preferences are significant predictors of one's own visitations (Chang and Sun 2011; Yang et al. 2013; Gao, Tang, and Liu 2012). To incorporate information from this modality, we first take the same three count vectors (non-normalized) from all of the friends of a user. For each location category, we run max-pooling to pick the highest visitation count among all friends. We also experimented with mean-pooling, but max-pooling gave better results. We do max-min normalization on each of the max-pooled vectors separately, to give us three aggregated visitation vectors from the friends, $Z_F^{tot}, Z_F^d, Z_F^h \in \mathbb{R}^D$, analogous to the user's own visitation frequency features. Concatenating the vectors gives us the friends' aggregated visitation feature vector, $Z_F = Z_F^{tot} \oplus Z_F^d \oplus Z_F^h$, where $Z_F \in \mathbb{R}^{3D}$.

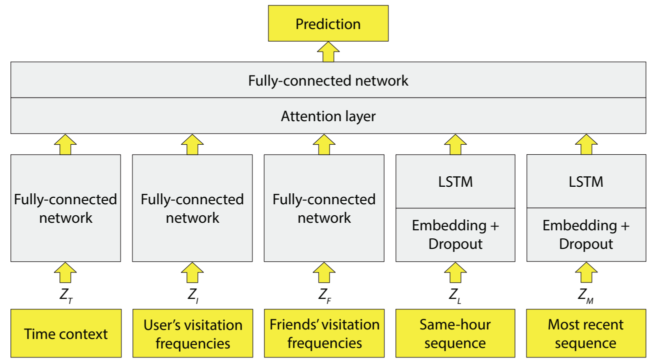

Figure 3: The LocHabits model architecture.

Most-recent category sequence. We take a sequence of $n_{seq,M}$ most-recent, time-stamped location categories. Each entry in the sequence is represented by concatenating three one-hot vectors denoting the location category ($\mathbb{R}^D$), day of the week ($\mathbb{R}^{D_d}$), and hour in the day ($\mathbb{R}^{D_h}$) respectively. This gives us the most-recent sequence feature representation, $Z_M \in \mathbb{R}^{n_{seq,M} \times (D + D_d + D_h)}$.

Same-hour category sequence. We also use same-hour category sequences to capture cyclical or long-term periodicity in the data (Chowdhury et al. 2021). We take a sequence of the user's $n_{seq,L}$ same-hour categories in the days before the query. We use zero-padding if a data point has fewer same-hour locations in previous days than $n_{seq,L}$. The locations are represented the same way as the most-recent location sequence, giving us the same-hour sequence feature representation, $Z_L \in \mathbb{R}^{n_{seq,L} \times (D + D_d + D_t)}$.

Model Architecture

The model architecture is shown in Figure 3. The five feature vectors go through their individual processing modules and get fed to an attention layer, A. All of the modeling choices were made through experimentation.

Time context module. The time-context feature Z_T is passed through a fully-connected network combined with a non-linear component. This helps create non-linear projections of the time context feature to capture the interactions between the day of the week and hour in the day features. This results in $Z_T^A = \mathcal{F}_T(Z_T)$, where $\mathcal{F}(\cdot)$ denotes a fully-connected network with dropout and activation, $Z_T^A \in \mathbb{R}^{D_A}$, and D_A is the hidden dimension of the attention layer.

Visitation feature modules of users and their friends. The visitation frequency feature vector of the user and

Dataset	Model	NDCG@1	NDCG@5	Recall@5	RR@5	Coverage Error	Macro-F1
Sn-CA-Jun	MCI-DNN	46.86%	68.74%	87.43%	62.49%	2.99	0.28
	SERM	46.47%	68.64%	**87.50%**	62.33%	**2.99**	0.28
	LocHabits	**50.68%**	**70.31%**	87.05%	**64.72%**	3.17	**0.34**
Sn-CA-Aug	MCI-DNN	48.34%	69.84%	88.00%	63.75%	2.93	0.30
	SERM	48.39%	69.96%	**88.15%**	63.86%	**2.91**	0.31
	LocHabits	**54.86%**	**72.67%**	87.62%	**67.66%**	3.06	**0.37**
Sn-NY-Aug	MCI-DNN	50.10%	70.06%	87.02%	64.39%	**3.05**	0.30
	SERM	49.46%	69.88%	**87.20%**	64.08%	3.06	0.30
	LocHabits	**51.89%**	**70.62%**	86.46%	**65.32%**	3.34	**0.33**
Sn-TX-Aug	MCI-DNN	49.74%	70.68%	87.95%	64.87%	2.93	0.26
	SERM	49.91%	70.82%	**88.03%**	65.03%	**2.92**	0.27
	LocHabits	**56.63%**	**73.69%**	87.79%	**68.95%**	2.99	**0.31**
Fsq-US-Apr	MCI-DNN	30.06%	53.46%	74.40%	46.50%	5.22	0.10
	SERM	31.56%	54.86%	75.58%	47.97%	4.99	0.11
	LocHabits	**37.41%**	**58.65%**	**77.33%**	**52.43%**	**4.68**	**0.16**
Fsq-JP-Apr	MCI-DNN	45.51%	66.29%	83.87%	60.40%	4.19	0.10
	SERM	47.54%	68.17%	85.28%	62.42%	3.92	0.11
	LocHabits	**54.49%**	**70.64%**	**88.27%**	**64.77%**	**3.14**	**0.15**
Fsq-BR-Apr	MCI-DNN	29.38%	52.67%	73.77%	45.67%	5.37	0.10
	SERM	31.45%	54.65%	75.40%	47.76%	5.06	0.11
	LocHabits	**41.08%**	**61.42%**	**79.31%**	**55.47%**	**4.19**	**0.18**

Table 3: Test-set performances of various models across datasets. The best performances in each dataset are highlighted in bold.

the aggregated visitation feature vector of their friends go through two separate modules. Each module consists of a fully connected network with dropout and activation layers, to capture the non-linear interactions among the overall, day-based, and hour-based features. This gives us $Z_I^A = \mathcal{F}_I(Z_I)$ and $Z_F^A = \mathcal{F}_F(Z_F)$, where Z_I^A, $Z_F^A \in \mathbb{R}^{D_A}$.

Sequence modules. The most recent and same-hour sequences go through two similar recurrent modules. In each module, the inputs are passed through embedding (Emb) and dropout (Drp) layers, to give $Z'_M = Drp(Emb(Z_M))$ and $Z'_L = Drp(Emb(Z_L))$ respectively for the most-recent and same-hour sequences. Here, $Z'_M, Z'_L \in \mathbb{R}^{n_{seq,M} \times D_e}$, where D_e denotes the embedding dimension. The embedding layer helps to deal with sparse features through linear transformations. The embedded representations are then passed through LSTM blocks to get $Z_M^A = LSTM_{out}(Z'_M)$ and $Z_L^A = LSTM_{out}(Z'_L)$, where $LSTM_{out}$ denotes the final hidden state output, and $Z_M^A, Z_L^A \in \mathbb{R}^{D_A}$. These outputs capture summary representations of the sequences.

Attention layer. The outputs of the five modules, Z_T^A, Z_I^A, Z_F^A, Z_M^A, and Z_L^A, are fed to an attention layer. A naive option is to concatenate the five outputs together, which gives every module equal importance. Instead, we use the attention mechanism, which attenuates or prioritizes each module's outputs adaptively for each prediction data point and helps improve performance. This also allows for a richer understanding of each module's impact on the prediction. The attention layer generates a soft-attended attention vector for each branch $z \in \{Z_T^A, Z_I^A, Z_F^A, Z_M^A, Z_L^A\}$ calculated as $\alpha_z = \exp(\phi(z))/\sum_z \exp(\phi(z))$, where $\phi(\cdot)$ is a mapping function implemented as a fully connected layer. Then, a fused embedding vector is computed as $Z_A = \sum_z \alpha_z z$.

Final prediction. We pass the fused embedding vector through another fully-connected network, before applying softmax to obtain a final prediction, $\mathrm{output} = softmax(\mathcal{F}_{out}(Z_A))$.

Results and Discussion

We train and test the LocHabits model using the data of 70k selected users (and their friends) from the Sn-CA-Jun dataset. This trained model is additionally tested on 21k selected users from each of the other Snapchat dataset versions. For the FourSquare datasets, we use the data from all of the users in the Fs-US-Apr dataset for training and testing a model. This trained model is additionally tested on the Fs-JP-Apr and Fs-BR-Apr datasets.

Comparison Baselines

Alongside the approaches mentioned in the Exploratory analysis section, we consider two high-performing neural models, MCI-DNN (Liao et al. 2018b) and SERM (Yao et al. 2017) as baselines. These models originally focused on predicting POIs rather than location categories, so we adapt the models to suit our problem scope. This gives us the following baselines:

- **Order-z Markov Model:** Exploratory analysis; Table 2.
- **Adapted MCI-DNN (Liao et al. 2018b):** The MCI-DNN model uses a ⟨location ID, context, timestamp⟩ POI tuple, and predicts the next location ID. To adapt the model to our scope, we denote each entry as a ⟨category ID, timestamp⟩ tuple, and predict the next category ID.

Model	NDCG@1	NDCG@5
LocHabits	50.68%	70.31%
- User's visitation frequencies	**49.00%**	**69.22%**
- Most recent sequence	**49.91%**	**69.47%**
- Time context	50.47%	70.16%
- Friends' visitation frequencies	50.57%	70.23%
- Same-hour sequence	50.59%	70.25%

Table 4: Impacts of different modules. Here '-' denotes the removal of the corresponding module.

Features used	NDCG@1	NDCG@5
Time context+User's visitation freq.		
+ Friends' visitation freq.	49.83%	69.44%
+ Same-hour sequence	49.91%	69.47%
+ Most recent sequence	50.68%	70.31%

Table 5: Incremental benefits of adding sequential data.

- **Adapted SERM (Yao et al. 2017):** Instead of describing location entries as a ⟨timestamp, location ID, activity description⟩ tuple, we denote them as ⟨category ID, timestamp⟩. We predict the next category ID instead of the next location ID. Furthermore, the SERM model represents users using one-hot encoding, which prohibits the model from readily generalizing to unseen users. To ensure generalizability and fair comparison, we represent each user by their overall visitation frequency vector.

Prediction Performance

The test-set performances are presented in Table 3. The LocHabits model outperforms the baseline results in all but two metrics (Recall@5 and CE) in the Snapchat dataset versions, and in all of the metrics in the FourSquare dataset versions. In particular, the LocHabits model outperforms the baselines by 14.94% on average in terms of getting the best guesses correct (NDCG@1)—which matters greatly in the content personalization setting. The model's ranking ordering is also better than the comparisons in the NDCG@5 and RR@5 metrics in all datasets. Notably, the baseline models predict the next venues without explicitly prioritizing the next-hour ones and do not incorporate features from social graph data and long-term patterns (i.e., same-hour sequence). These can potentially explain their lack of NDCG@1 performance compared to LocHabits.

The models trained on the `Sn-CA-Jun` dataset sustain their performances in the test sets of datasets from different areas (`Sn-NY-Aug`, and `Sn-TX-Aug`) even at a later time (2 months later; `Sn-CA-Aug`, `Sn-NY-Aug`, and `Sn-TX-Aug`), without requiring re-training. The same trend is seen for the FourSquare datasets. This shows that the models are generalizable to out-of-sample users from other places and times.

Impact of Different Modules

Ablation study. To study the role of different features, we remove one module at a time and retrain the reduced models. In Table 4, we report the reduced model performances on the test set of the `Sn-CA-Jun` dataset. The first row in Table 4 shows the performance of the full model. From there, removing each of the modules reduces the performance, implying that each of those makes a notable contribution to the final performance. The highest reductions in performance are observed when the user's visitation frequencies and the most recent location sequences are removed, respectively marking the top two most informative feature modules.

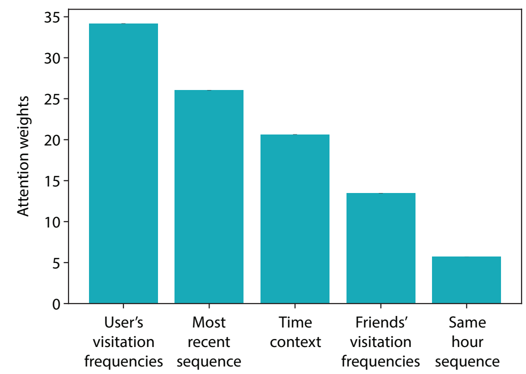

Figure 4: Attention weights of various feature modules.

Although we scoped our task and built the LocHabits model around predicting the next-hour category, it is possible to extend the scope and framework to predict categories at any arbitrary future hour. In particular, the modules of time context, user's visitation frequencies, and friends' aggregated visitation frequencies are suited to make predictions at arbitrary query times using the temporal regularities they capture. In the scope of our task, these three modules lead to a prediction performance of NDCG@1 = 49.83% for the `Sn-CA-Jun` dataset (first row in Table 5). The same-hour sequence module (many-to-one LSTM) can be used to influence the predictions for as far as 24 hours into the future. In our task, adding this module (row 2 in Table 5) leads to improved performance of 49.91%. Finally, the most recent sequence module (also a many-to-one LSTM) is limited to influencing the results only for the next hour. In our task, adding this module (i.e., full model) gives a performance of 50.68% (row 3 in Table 5). If the same-hour sequence and most recent sequence modules are to make predictions at any arbitrary future hour—which may compromise performance due to looking too far ahead in time—these two modules will need to be replaced by suitable many-to-many/multiple many-to-one recurrent blocks.

Attention weights. As described earlier, the attention layer attenuates or prioritizes each feature for each prediction data point adaptively. Thus, on the test set, we can observe which features are getting more attention, informing us of their relative importance. The attention weights for the five sets of features are shown in Figure 4, as collected from the test set of the `Sn-CA-Jun` dataset. As can be seen, the users' visitation frequencies and the most recent location sequences are once again the most informative features, cor-

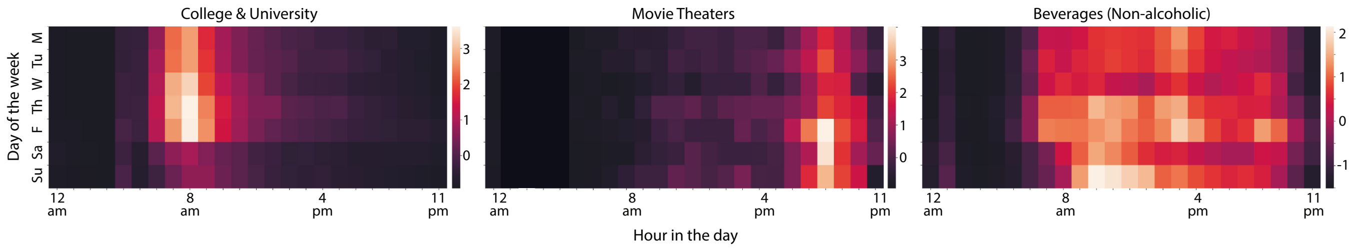

Figure 5: Users' visitation heatmaps shown for three example location categories from the `Sn-NY-Aug` dataset. The entries are z-normalized, and brighter colors denote higher visitation counts. The users typically visit the 'College & University' category in the early mornings of weekdays, and 'Movie theaters' on the evenings of Friday to Sunday. The visitations to the 'Beverages (non-alcoholic)' category are spread out during the daytime throughout the week.

roborating the insights from the ablation study. All of the feature sets contribute with weights significantly above 0%.

Cold start scenario. The model does not learn any user-specific parameters and can be readily used in cold-start scenarios where there is no user data available to construct the historical feature vectors. In that case, using only the time-context module achieves NDCG@1=26.97%, capturing the collective trends associated with the time-context features. The performance is on par with the collective preference-based results shown in the exploratory analyses.

Qualitative Explorations

- **Contrasting local and global patterns.** Visiting the 'College & University' category on weekdays is a predictable behavior for users across various states. Whereas, visiting 'Harbor & Marina' is a behavior mostly found in places that have better access to water lines. Going to the coffee shop appears to be a visitation behavior that can look different across users. These properties of data distribution are reflected in the neural model's prediction performances: the more regularity there is in the data, the more predictable a category is. For example, the 'College & University' category comes with a decent predictability (per-class F1 score $\sim$ 55%), while the 'Beverages (non-alcoholic)' category, which includes coffee shops, has much lower predictability ($\sim$ 11%) (see Figure 5).
- **Capturing temporal effects.** In June, there are fewer data points for 'College & University', due to it being summertime in the US when classes are off. In August, we notice substantially increased entries in this category. This is reflected in the category's prediction performance, which is substantially higher in August (per-class F1 score $\sim$ 65%) than in June ($\sim$ 55%). Analyzing the temporal changes in the per-class F1 scores can thus function as a potential marker for regularity in people's category visitations.

Limitations and Future Work

We used graph features from a single hop of neighbors and gave each friend an equal weight. It is worth experimenting with multiple hops of message propagation in the social graph. The social influence on people's future location categories can vary based on friendship strength. In future work, we will incorporate weighted graphs to selectively focus on the most informative/influential friends. Taking only the top 50 most frequent categories helped save computational resources while catering to the vast majority (99.29%) of the users' entries. However, this choice limited the model's ability to predict when people might visit extremely rare categories. We used one-hot encoded vectors for representing location categories in our work. In future work, we will experiment with embeddings that capture richer semantic characteristics of the categories.

Conclusion

We predicted people's location categories in the next hour with a privacy-first approach. In contrast to most approaches, we avoided spatial coordinates along with other user-identifiable sources and only used location *categories*, timestamps, and social graph data. We then curated feature sets informed by previous works and our own explorations. We crafted a neural model that is generalizable and works well in cold-start situations. We found the user's visitation frequencies and most recent sequences to be the most predictive features. Our work can enable rich personalization of content delivery on online social platforms.

Ethics Statement

Any experiment dealing with data as sensitive as ours (e.g., location) needs to operate ethically and securely. Our approach actively aims to minimize risks of misuse and intrusion by avoiding user-identifiable data, such as demographic identities and spatial coordinates. Thus, our model may be preferable in highly sensitive settings. The datasets were anonymized before analysis. All experiments were conducted in Snapchat's internal secure storage systems, and no data was stored outside Snapchat's ecosystem. Thus, we do not foresee strong ethical concerns induced by our work.

References

Aarts, H.; and Dijksterhuis, A. 2000. The Automatic Activation of Goal-directed Behaviour: The Case of Travel Habit. *Journal of Environmental Psychology*, 20(1): 75–82.

Berkovsky, S.; Kuflik, T.; and Ricci, F. 2008. Mediation of user models for enhanced personalization in recommender systems. *User Modeling and User-Adapted Interaction*, 18(3): 245–286.

Burbey, I. 2011. *Predicting future locations and arrival times of individuals*. Ph.D. thesis, Virginia Tech.

Cao, J.; Xu, S.; Zhu, X.; Lv, R.; and Liu, B. 2018. Effective Fine-grained Location Prediction based on User Check-in Pattern in LBSNs. *Journal of Network and Computer Applications*, 108: 64–75.

Chang, B.; Park, Y.; Park, D.; Kim, S.; and Kang, J. 2018. Content-Aware Hierarchical Point-of-Interest Embedding Model for Successive POI Recommendation. In *IJCAI*, 3301–3307.

Chang, J.; and Sun, E. 2011. Location3: How Users Share and Respond to Location-Based Data on Social. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1): 74–80.

Chen, M.; Liu, Y.; and Yu, X. 2014. Nlpmm: A next location predictor with Markov modeling. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 186–197. Springer.

Chen, Z.; Kelty, S.; Evsukoff, A. G.; Welles, B. F.; Bagrow, J.; Menezes, R.; and Ghoshal, G. 2022. Contrasting social and non-social sources of predictability in human mobility. *Nature Communications*, 13(1): 1–9.

Cheng, C.; Yang, H.; Lyu, M. R.; and King, I. 2013. Where you like to go next: Successive point-of-interest recommendation. In *Twenty-Third International Joint Conference on Artificial Intelligence*.

Cho, E.; Myers, S. A.; and Leskovec, J. 2011. Friendship and mobility: User movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1082–1090.

Chowdhury, F. A.; Liu, Y.; Saha, K.; Vincent, N.; Neves, L.; Shah, N.; and Bos, M. W. 2021. CEAM: The Effectiveness of Cyclic and Ephemeral Attention Models of User Behavior on Social Platforms. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 117–128.

Christoforidis, G.; Kefalas, P.; Papadopoulos, A.; and Manolopoulos, Y. 2018. Recommendation of Points-of-Interest using graph embeddings. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, 31–40. IEEE.

Christoforidis, G.; Kefalas, P.; Papadopoulos, A. N.; and Manolopoulos, Y. 2021. RELINE: Point-of-Interest recommendations using multiple network embeddings. *Knowledge and Information Systems*, 63(4): 791–817.

Cui, Q.; Zhang, C.; Zhang, Y.; Wang, J.; and Cai, M. 2021. ST-PIL: Spatial-Temporal Periodic Interest Learning for Next Point-of-Interest Recommendation. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2960–2964.

Cui, R.; Agrawal, G.; and Ramnath, R. 2020. Tweets can tell: Activity recognition using hybrid gated recurrent neural networks. *Social Network Analysis and Mining*, 10(1): 1–15.

De Montjoye, Y.-A.; Hidalgo, C. A.; Verleysen, M.; and Blondel, V. D. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3(1): 1–5.

Doan, K. D.; Yang, G.; and Reddy, C. K. 2019. An Attentive Spatio-Temporal Neural Model for Successive Point of Interest Recommendation. In *PAKDD (3)*, 346–358.

Elmi, S.; Benouaret, K.; and Tan, K.-L. 2021. Social and Spatio-Temporal Learning for Contextualized Next Points-of-Interest Prediction. In *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, 322–329. IEEE.

Gao, H.; Tang, J.; and Liu, H. 2012. Exploring Social-Historical Ties on Location-based Social Networks. In *Sixth International AAAI Conference on Web and Social Media*.

Gao, Q.; Zhou, F.; Trajcevski, G.; Zhang, K.; Zhong, T.; and Zhang, F. 2019. Predicting Human Mobility via Variational Attention. In *The World Wide Web Conference*, 2750–2756.

Guo, Q.; Sun, Z.; Zhang, J.; and Theng, Y.-L. 2020. An attentional recurrent neural network for personalized next location recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 83–90.

He, J.; Li, X.; and Liao, L. 2017. Category-aware Next Point-of-Interest Recommendation via Listwise Bayesian Personalized Ranking. In *IJCAI*, volume 17, 1837–1843.

Islam, M. A.; Mohammad, M. M.; Das, S. S. S.; and Ali, M. E. 2022. A Survey on Deep Learning based Point-of-Interest (POI) Recommendations. *Neurocomputing*, 472: 306–325.

Kefalas, P.; and Manolopoulos, Y. 2017. A time-aware spatio-textual recommender system. *Expert Systems with Applications*, 78: 396–406.

Kurashima, T.; Althoff, T.; and Leskovec, J. 2018. Modeling interdependent and periodic real-world action sequences. In *Proceedings of the 2018 World Wide Web Conference*, 803–812.

Li, M.; Zheng, W.; Xiao, Y.; Zhu, K.; and Huang, W. 2021. Exploring Temporal and Spatial Features for Next POI Recommendation in LBSNs. *IEEE Access*, 9: 35997–36007.

Li, R.; Shen, Y.; and Zhu, Y. 2018. Next Point-of-Interest recommendation with temporal and multi-level context attention. In *2018 IEEE International Conference on Data Mining (ICDM)*, 1110–1115. IEEE.

Lian, D.; and Xie, X. 2011. Collaborative Activity Recognition via Check-in History. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, 45–48.

Liao, D.; Liu, W.; Zhong, Y.; Li, J.; and Wang, G. 2018a. Predicting Activity and Location with Multi-task Context Aware Recurrent Neural Network. In *IJCAI*, 3435–3441.

Liao, J.; Liu, T.; Liu, M.; Wang, J.; Wang, Y.; and Sun, H. 2018b. Multi-context integrated deep neural network model for next location prediction. *IEEE Access*, 6: 21980–21990.

Liu, Q.; Wu, S.; Wang, L.; and Tan, T. 2016. Predicting the next location: A recurrent model with spatial and temporal contexts. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Liu, T.; Liao, J.; Wang, Y.; Wang, J.; and Qi, Q. 2019. Collaborative tensor–topic factorization model for personalized activity recommendation. *Multimedia Tools and Applications*, 78(12): 16923–16943.

Navet, N.; and Chen, S.-H. 2008. On Predictability and Profitability: Would gp Induced Trading Rules be Sensitive to the Observed Entropy of Time Series? In *Natural Computing in Computational Finance*, 197–210. Springer.

Nawshin, S.; Mukta, M. S. H.; Ali, M. E.; and Islam, A. N. 2020. Modeling Weather-Aware Prediction of User Activities and Future Visits. *IEEE Access*, 8: 105127–105138.

Noulas, A.; Scellato, S.; Lathia, N.; and Mascolo, C. 2012. Mining User Mobility Features for Next Place Prediction in Location-based Services. In *2012 IEEE 12th International Conference on Data Mining*, 1038–1043. IEEE.

Ochiai, K.; Fukazawa, Y.; Yamada, W.; Manabe, H.; and Matsuo, Y. 2020. Gravity of location-based service: Analyzing the effects for mobility pattern and location prediction. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, 476–487.

Pyrgelis, A.; Troncoso, C.; and De Cristofaro, E. 2017. Knock knock, who's there? Membership inference on aggregate location data. *arXiv preprint arXiv:1708.06145*.

Qiao, S.; Han, N.; Zhou, J.; Li, R.-H.; Jin, C.; and Gutierrez, L. A. 2018. SocialMix: A familiarity-based and preference-aware location suggestion approach. *Engineering Applications of Artificial Intelligence*, 68: 192–204.

Rahimiaghdam, S.; Karagoz, P.; and Mutlu, A. 2016. Personalized time-aware outdoor activity recommendation system. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 1121–1126.

Rahmani, H. A.; Aliannejadi, M.; Baratchi, M.; and Crestani, F. 2020. Joint geographical and temporal modeling based on matrix factorization for Point-of-Interest recommendation. *Advances in Information Retrieval*, 12035: 205.

Schein, A. I.; Popescul, A.; Ungar, L. H.; and Pennock, D. M. 2002. Methods and Metrics for Cold-start Recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 253–260.

Song, C.; Qu, Z.; Blumm, N.; and Barabási, A.-L. 2010. Limits of Predictability in Human Mobility. *Science*, 327(5968): 1018–1021.

Sun, K.; Qian, T.; Chen, T.; Liang, Y.; Nguyen, Q. V. H.; and Yin, H. 2020. Where to go next: Modeling long-and short-term user preferences for point-of-interest recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 214–221.

To, D.; Si, D.; and Chen, Y. 2019. Traveler's Next Activity Predication with Location-Based Social Network Data. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Prediction of Human Mobility*, 15–23.

Wang, S.; Li, A.; Xie, S.; Li, W.; Wang, B.; Yao, S.; and Asif, M. 2021. A Spatial-Temporal Self-Attention Network (STSAN) for Location Prediction. *Complexity*, 2021.

Wu, Y.; Li, K.; Zhao, G.; and Xueming, Q. 2022. Personalized long-and short-term preference learning for next POI recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 34(4): 1944–1957.

Xie, M.; Yin, H.; Wang, H.; Xu, F.; Chen, W.; and Wang, S. 2016. Learning graph-based POI embedding for location-based recommendation. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 15–24.

Yang, C.; Sun, M.; Zhao, W. X.; Liu, Z.; and Chang, E. Y. 2017. A neural network approach to jointly modeling social networks and mobile trajectories. *ACM Transactions on Information Systems (TOIS)*, 35(4): 1–28.

Yang, D.; Fankhauser, B.; Rosso, P.; and Cudre-Mauroux, P. 2020. Location Prediction over Sparse User Mobility Traces Using RNNs: Flashback in Hidden States! In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2184–2190.

Yang, D.; Qu, B.; Yang, J.; and Cudre-Mauroux, P. 2019. Revisiting User Mobility and Social Relationships in LBSNs: A Hypergraph Embedding Approach. In *The World Wide Web Conference*, 2147–2157.

Yang, D.; Zhang, D.; Yu, Z.; and Wang, Z. 2013. A Sentiment-enhanced Personalized Location Recommendation System. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, 119–128.

Yang, D.; Zhang, D.; Zheng, V. W.; and Yu, Z. 2014. Modeling User Activity Preference by Leveraging User Spatial Temporal Characteristics in LBSNs. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(1): 129–142.

Yao, D.; Zhang, C.; Huang, J.; and Bi, J. 2017. SERM: A recurrent model for next location prediction in semantic trajectories. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2411–2414.

Ye, J.; Zhu, Z.; and Cheng, H. 2013. What's Your Next Move: User Activity Prediction in Location-based Social Networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, 171–179. SIAM.

Yu, F.; Cui, L.; Guo, W.; Lu, X.; Li, Q.; and Lu, H. 2020. A category-aware deep model for successive poi recommendation on sparse check-in data. In *Proceedings of the Web Conference 2020*, 1264–1274.

Zhang, D.; Wang, D.; Zheng, H.; Mu, X.; Li, Q.; and Zhang, Y. 2017. Large-scale point-of-interest category prediction using natural language processing models. In *2017 IEEE International Conference on Big Data (Big Data)*, 1027–1032. IEEE.