# Activity Sensor: Check-In Usage Mining for Local Recommendation

JITAO SANG, Chinese Academy of Sciences
TAO MEI, Microsoft Research Asia
CHANGSHENG XU, Chinese Academy of Sciences



While on the go, people are using their phones as a personal concierge discovering what is around and deciding what to do. Mobile phone has become a recommendation terminal customized for individuals—capable of recommending activities and simplifying the accomplishment of related tasks. In this article, we conduct usage mining on the check-in data, with summarized statistics identifying the local recommendation challenges of huge solution space, sparse available data, and complicated user intent, and discovered observations to motivate the hierarchical, contextual, and sequential solution. We present a point-of-interest (POI) category-transition–based approach, with a goal of estimating the visiting probability of a series of successive POIs conditioned on current user context and sensor context. A mobile local recommendation demo application is deployed. The objective and subjective evaluations validate the effectiveness in providing mobile users both accurate recommendation and favorable user experience.

Categories and Subject Descriptors: H.4 [**Information Systems Applications**]: Miscellaneous

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Location-based service, local recommendation, check-in, usage mining

**ACM Reference Format:**
Jitao Sang, Tao Mei, and Changsheng Xu. 2015. Activity sensor: Check-in usage mining for local recommendation. ACM Trans. Intell. Syst. Technol. 6, 3, Article 41 (April 2015), 24 pages.
DOI: http://dx.doi.org/10.1145/2700468

## 1. INTRODUCTION

By September 2013, 56% of people in the world owned a smart mobile device, among which more than half used mobile as their primary Internet access portal, making the number of mobile Internet users reach 1.36 billion. Various mobile Internet applications provide massive opportunities to facilitate users, such as pervasive Internet, music listening and mobile reading at fragment time, and mobile recommendation for offline transactions. The ubiquitous smart mobile devices have revolutionized the way people surf information and make decisions. People are using their mobile phones to accomplish tasks anytime and anywhere. Statistics shows that 90% of users who query local information on mobile devices conduct successive activities, such as call for further inquiry or finish with transactions. During the shopping festival "Black Friday" in 2012, 24% of the overall transactions came from mobile terminals. It is therefore

This work is supported in part by the National Basic Research Program of China (2012CB316304), National Natural Science Foundation of China (61225009, 61303176, 61272256), and Beijing Natural Science Foundation (4131004).
Q1 Authors' addresses: J. Sang and C. Xu, National Lab of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences; T. Mei, Microsoft Research Asia.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2015 ACM 2157-6904/2015/04-ART41 $15.00
DOI: http://dx.doi.org/10.1145/2700468

important to understand user intent on the go toward intelligent local recommendation services.

To understand user intent in the context of local recommendation services, we aim to predict which point of interest (POI), such as a shopping mall or a western restaurant, a user is likely to visit next. Two problems need to be involved toward this goal: how to explore usage patterns from the large-scale, noisy, and heterogeneous user mobile data, and how to employ the discovered usage pattern to perform local recommendation customized to contextual and personalized information needs. In this work, we resort to exploit the user check-in data to motivate solutions, which is initiated by Foursquare to allow users self-reported position to a physical place and share their locations with friends. Although traditional records of mobile cellular data [Gonzalez et al. 2008], query log [Zhuang et al. 2011], and click-through history [Lane et al. 2010] have been widely used for mobile behavior modeling, check-in data differentiates in three ways. First, check-ins naturally bridge user intent and the behaviors in the physical world. The place where a user chooses to check in usually corresponds to a specific POI. Therefore, the discovered patterns from check-in usage mining are easily incorporated into POI-oriented local recommendation applications. Second, check-in provides rich context of user behavior and status. A check-in POI serves as a good indicator of the user's current status and the activity being conducted, which helps to refine the complicated user intent. For example, a user's check-in at a plaza implies her shopping status/activity and suggests the intent for having tea or eating. Third, check-in data is publicly available. Different from mobile browsing log data, in which users have no intent to broadcast, in check-in services, users voluntarily share their status, making it suitable for large-scale research use. These three characteristics make check-in data a promising source for mobile usage mining and local recommendation.

We conduct a comprehensive check-in usage mining on a real-world check-in dataset. The statistics contribute to identification of research challenges, and the observations contribute to inspire the latter local recommendation solution. In particular, local recommendation is challenging in three ways. First, the solution space for local recommendation is huge. Besides the notorious positioning error, even when we know exactly where users locate, there exist large number of POI alternatives for choice. According to the statistics, the average distance that the user is likely to travel for the next action is 0.33km. Considering the real POI frequency in our collected Shanghai dataset, the number of candidate POIs equals 108 if 0.33km is the maximum distance for the next action. Second, the available user data for usage mining is sparse. Compared to the scale of POI, a user usually has a very limited number of publicly shared behavior data. According to our statistics on the 5-month check-in dataset, the average number of check-in records per month is only 11. The challenge lies in how to exploit these limited check-in data for accurate user modeling and practical recommendation. Last, but maybe the most important, user intent is complicated. In many cases, the user does not have a very explicit intent in mind but just wants some suggestions for the desired activity. Our statistics reveal that users are likely to conduct several POI actions in short period, which indicates that user activity usually consists of a package of sequential actions. For example, a typical "nightout" activity will include several POI actions of having dinner, going to Karaoke, and so on.

To address the preceding challenges, inspired by the observations from check-in usage mining, we propose our solution for hierarchical, contextual, and sequential local recommendation. First, from the *hierarchical* view, to conduct successful local POI recommendation, it is of significance to investigate the category of POI. In this article, we refer to a POI category as a set of POIs with the same functionality, such as "shopping mall" and "restaurant." It has been recognized that user activity intent usually corresponds to a certain POI category instead of a specific POI [Zhuang et al. 2011]. Our observation from check-in usage mining shows an obvious transition pattern in

Fig. 1. The interface of Activity Sensor developed on a Windows Phone device. A user named "Emily" checks in to a shopping mall, providing user information and context (a), gets recommendation of sequential POI categories (e.g., cafe→restaurant) (b), and the gets corresponding packages of sequential POIs when selecting a specific POI category sequence (c).

the POI category level, based on which we present a hierarchical Markov solution by first recommending a POI category to clarify the user intent and thus reduce the number of POI candidates (Figure 1). Second, from the *contextual* view, the long-term preference discovered from check-in history is insufficient for local recommendation for two reasons: user data sparsity and the dynamic characteristic of user intent. In this work, we integrate long-term preference with short-term intent by considering user context (e.g., current check-in record) as well as sensor contexts (e.g., time, location). Although time and location connect users with similar sensor context for efficient collaborative filtering (CF), the current check-in record provides an important indication to understand the on-site intent. Third, from the *sequential* view, the reasonability of POI recommendation is easily validated from people's tendency to conduct successive actions, where each action can correspond to one single check-in POI. To simulate the process of planning activities in real life, instead of recommending a single POI as most existing works did, we propose recommendation of a package of sequential POIs. For example, the recommendation of "a tea break in some nearby cafe" and then "a decent western restaurant" to a lady who currently checks in to a shopping mall would be a nice plan for her. In the broader context, sensing and predicting users' next possible check-in POIs can facilitate users in exploring and planning activities, which will go beyond check-in services and extend the dimension of current location-based services.

To summarize, this work makes the following contributions:

—We conduct an in-depth check-in usage mining. The resultant statistics and observations serve as both proof to identify research challenges and a foundation to inspire practical solutions. It also provides good knowledge for future applications based on check-in data.
—To alleviate the user data sparsity, we consider both sensor and user context to develop a contextual recommendation algorithm. Regarding the large number of POI candidates, a hierarchical scheme is designed for coarse-to-fine POI recommendation.
—We address the local recommendation problem by moving one step to recommend a series of sequential POIs. This is more natural for humans to plan their activities in real life.
—We develop a real-world application based on the proposed approach on Windows Phone 7 devices, with both objective and subjective experimental results for evaluation.

## 2. RELATED WORK

### 2.1. Usage Mining

Usage mining refers to the discovery of useful knowledge from various user behavior logs. It has been investigated in the context of web mining, which focuses on finding patterns of navigational behavior from users visiting a Web site [Cooley et al. 1999]. Web usage mining has served as a necessary task to provide Web administrators with meaningful information about users and usage patterns for improving quality of Web information and service performance; for example, it is customized to meet user preferences either in the presentation of information or in relevance to the content that best fits the user [Wang et al. 2005; Lappas 2011].

The widespread use of mobile devices has offered the opportunity to gain insights on user mobile data and useful behavior patterns. Mobile usage mining has been explored on cellular data, mobile query logs, and click-through histories. In Cranshaw et al. [2010], the authors examined the traces of users by recording their cell phone traces and found the shared patterns between the mobility patterns and social networks. In Lane et al. [2010], a the Hapori system builds successful user behavioral models by mining the similarity between users' mobile click-through history. In Zhuang et al. [2011], mobile query data is exploited to discover the correlations between user mobile activity and sensor contexts.

Recently, check-in data have been investigated in various applications, including popular POI identification [Alves et al. 2011], user activity recognition [Lian and Xie 2011], friendship prediction [Sadilek et al. 2012], and user localization [Cho et al. 2011]. One research line of usage mining around check-in data has focused on basic geographical or temporal distribution [Ye et al. 2011; Noulas et al. 2011], as well as check-in frequency [Chang and Sun 2011]. Another line is to exploit the correlation between check-in action and other factors for solutions. Typical works include analyzing between geographical and social information for location prediction [Gao et al. 2012] and geographically clustering check-in records to associate user check-in actions with demographic attributes [Cheng et al. 2012]. However, the temporal and geographical relation between successive check-ins, although discussed in Noulas et al. [2011] and Cheng et al. [2013], has not been fully exploited. In this article, we discover a distinct category-transition pattern between successive check-ins, which serves as the basic premise for the hierarchical solution. Moreover, the POI repetition pattern is observed and discussed in the context of check-in data.

Note that various usage mining has been conducted toward recommendation tasks, especially on prediction of a user's next choice based on the current choice in sessions. The URL recommendation problem, which predicts next page requests, has been explored in Borges and Levene [2010] by analyzing the user navigation patterns to complete the trail with a given trail. In Davidson et al. [2010], video closeness is measured by co-visitation frequency, based on which an online video recommender system for YouTube is presented. The authors in Park et al. [2011] discussed a CF-based next-songs prediction method by analyzing the song-session similarities. In Rendle et al. [2010], the relationship of products between sessions are exploited from real online transaction data. The next-products recommendation solution is developed based on the discovered usage patterns. The key problem in next-choice recommendation lies in the relationship analysis with the observed action. The preceding works conduct relationship analysis on the item level. However, through our check-in data observation, we find significant POI category-transition patterns within sessions. Therefore, instead of focusing on item level, we propose a category-transition–based solution toward local recommendation.

Table I. Categorization of the Works on Local Recommendation According to Three Features

| Work | Hierarchical | Contextual | Sequential |
|---|---|---|---|
| Horozov et al. [2006], Ye et al. [2010] | | | |
| Zheng et al. [2010a], Lane et al. [2010], Kurashima et al. [2013] | | ✓ | |
| Zheng et al. [2010b], Liu and Xiong [2013] | ✓ | ✓ | |
| Chang et al. [2011], Wei et al. [2012] | | ✓ | ✓ |
| Activity Sensor | ✓ | ✓ | ✓ |

### 2.2. Local Recommendation

Local recommendation has received much attention in recent years due to the popularity of location-based social networks and the prevalence of various location-based services. The basic task for local recommendation is to consider users' location history as input and provide each user with a potentially interesting location to visit. In the general recommender system, works are organized based on their utilized methodologies into collaborative, content based, and hybrid categories [Adomavicius and Tuzhilin 2005]. Regarding the challenges that we summarized in the Introduction and the particular application scenario of local recommendation, in this section we attempt to categorize the representative local recommendation solutions according to the three functions they support: hierarchical, contextual, and sequential.

In Table I, we summarize the categorization of the existing works according to the three functions. Early works consider local recommendation as one special type of recommender system and resort to classic solutions without explicit modeling of sensor contexts. CF is widely utilized to provide personalized recommendations for locations based on similar users' historical locations or ratings. Horozov et al. [2006] uses the historical ratings from users living close to the user's current location, which significantly reduces the number of users in the user similarity matrix and thus reduces the computational cost of the recommendation. Ye et al. [2010] also employ a user-based CF model but suggest that solely using the ratings of a user's friends is more efficient and effective. Experiments show that a user's friends share more preferences than strangers.

A user's preferences and needs on the go change with the contexts [Sang et al. 2013]. It is important to integrate context into the location ranking for recommendation. A natural extension to CF on integrating context is tensor factorization, where context is treated as one of the factors. Zheng et al. [2010a] proposes a tensor factorization–based model to exploit user activity history for local recommendation. Symeonidis et al. [2013] extends the previous work by using the higher-order singular value decomposition (HOSVD) technique to perform contextual local recommendation. Instead of CF models, Kurashima et al. [2013] built a geographical topic model to understand user local behaviors in the same city. In addition to location context, richer context of time and weather are considered and exploited in Lane et al. [2010]. The Hapori system is developed to demonstrate the potential of integrating more contexts. Preceding contextual local recommendation systems exploit a user's history records to extract context-aware long-term preferences. However, the short-term needs, indicated by the user's current context, is largely ignored.

Local recommendation holds unique hierarchical property. On the one hand, locations span multiple scales [Bao et al. 2013]. Locations with different granularities form a hierarchy, where the geographical relation between different ties encodes important prior information. On the other hand, a user's intent toward the POIs is inherently function oriented. The category-POI hierarchical organization is both natural and efficient to a practical recommendation system. Recently, some works have exploited this

Table II. Examples of Check-In Records

| User ID | Time | Location | POI | Category |
|---|---|---|---|---|
| 2024771825 | 3/18/2011 9:44:20 PM | 31.229242, 121.446499 | Starbucks | tea & cafe |
| 1000246505 | 6/7/2011 8:01:33 PM | 32.043538, 118.785519 | Carrefour | supermarket |

property in their solutions. Zheng et al. [2010b] extend the hypertext induced topic search (HITS) model to extract interesting locations and experienced users by dividing the geographical space into a tree-based hierarchical graph (TBHG). Liu and Xiong [2013] propose a topic- and location-aware POI recommender system by exploiting associated POI category and context information.

Sequential local recommendation rises to address the complicated user intent and can satisfy more complex objectives. Chang et al. [2011] present a route recommendation solution that takes into account a user's own historically preferred road segments, mined from the user's historical GPS trajectories. Zheng and Xie [2011] present travel recommendation solutions by mining multiple users' GPS traces. Geotagged social media is also exploited to improve sequential local recommendation. Wei et al. [2012] propose the Route Inference framework based on Collective Knowledge (RICK) to construct popular routes from geotagged photos. Given a location sequence and a time span, RICK constructs the top-k routes by aggregating uncertain trajectories to facilitate trip recommendation. In this article, we propose a novel local recommendation system named "Activity Sensor" to integrate all three functions of contextual, hierarchical, and sequential: both sensor contexts (e.g., location and time) and user context (current check-in) are modeled to recommend users a series of successive POIs into a two-level category-POI organization.

## 3. CHECK-IN USAGE MINING

We collected a check-in dataset from Dianping,[1] the check-in Web site that cooperates with China's biggest microblog, Sina Weibo.[2] We conducted an in-depth analysis to investigate the characteristics of check-in records and transition patterns in successive check-in records, which further motivates the proposed recommendation solution.

The collected check-in records are located at Beijing and Shanghai from January 7 to June 11 in 2011. Each check-in record consists of user identification, check-in time, GPS latitude and longitude, check-in POI, and the POI category. Examples of check-in records are listed in Table II. We represent each check-in record as a five-dimensional tuple:

$$R := < U, T, L, O, C >,$$

where $U \in \mathcal{U}$ is the user ID,[3] $L$ and $T$ are the sensor contexts (i.e., location and time when $R$ is generated), $O \in \mathcal{O}$ is the check-in POI, and $C \in \mathcal{C}$ is the category of $O$. Besides the check-in records, we also collected the gender and residence information for each user, as well as the attributes of each POI, such as the name, category, GPS location, and address. In the following, we first summarize the basic statistics of the check-in dataset and then discuss several key observations.

### 3.1. Statistics

*3.1.1. Statistic I: #Check-In & POI Category.* The overall number of check-in records is $152,154$, issued by $2,342$ unique users located in Beijing and Shanghai. For the duplicate records where two successive check-in POIs are identical, we removed

[1]http://www.dianping.com/.
[2]http://weibo.com/.
[3]The user information is anonymous.

Fig. 2. The location distribution of user check-in to Beijing and Shanghai.

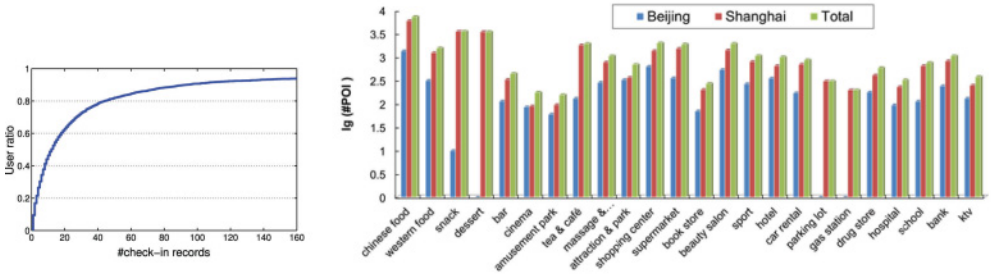


Fig. 3. Statistic I: CDF of #user check-in (left) and #POI in 24 categories (right).

the latter. After filtering, we obtained 128, 937 records in total. In Figure 2, we map the GPS locations of check-in records to Beijing and Shanghai, respectively. The depth of the red color on the map is proportional to the number of check-in records conducted at this area. We can see the check-in clusters easily identify the hot areas in each city, such as the central business district, the shopping areas, and dining areas. From these records, 32, 891 unique POIs are identified, resulting in 24 POI categories on public facilities, entertainment, life services, and so on. The number of POIs in Beijing and Shanghai for each category is shown on the right-hand sid of Figure 3. Regarding the #check-in statistics at user-level, the #check-in record for each user ranges from 1 to 3, 268. We draw the cumulative user ratio versus #check-in on the left-hand side of Figure 3, where we can see that more than 80% of users share no more than 50 check-in records. The average #check-in records for each user is only 55, and 11 in each month, which is very small compared to the overall #POI.

*3.1.2. Statistic II: Inter Check-In Time & Session Length.* Inter check-in time indicates the temporal interval between two successive check-in records for the same user [Noulas et al. 2011]. The left-hand side of Figure 4 shows the distribution of inter check-in time. It can be seen that the inter check-in time spans a long range, from 10 minutes to more than 2 hours. We notice that considerable numbers of successive check-ins are submitted within a very short period. Considering that it is unusual for one to conduct different actions in a very short time, most of these check-ins are for fun or for the sake of benefit and thus do not accord with real user behaviors [Lian and Xie 2011]. Therefore, we exclude the successive records with inter check-in time less than 10 minutes for further usage mining. On the other end, if inter check-in time between two records is very long (e.g., 2 hours), it is very likely that the latter check-in action has little relation with the former ones.

In this context, we define "session" to represent a series of related consistent check-in actions, which is constructed by those records with inter check-in time in

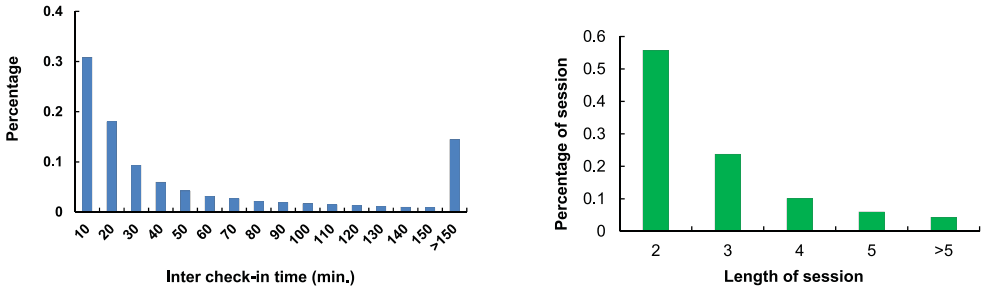


Fig. 4. Statistic II: Inter check-in time distribution (left) and session length distribution (right).

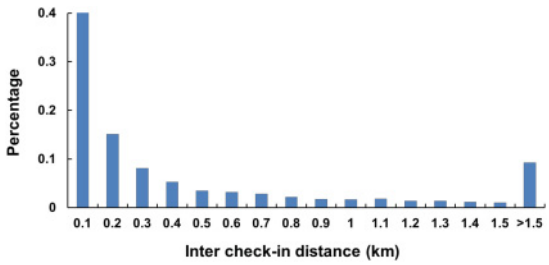


Fig. 5. Statistic III: Inter check-in distance distribution.

[10 min, 120 min]. Therefore, each session can be represented as

$$S := < U, T_1, L_1, O_1, C_1, \ldots, T_{N_S}, L_{N_S}, O_{N_S}, C_{N_S} >,$$

where $N_s$ is the length of session and $T_k$, $L_k$, $O_k$, and $C_k$ denote the time, location, POI name, and POI category of the $k$-th record in the session $S$, respectively. As a result, 13,078 sessions are extracted. The distribution of session lengths is shown on the right-hand side of Figure 4. We can see that around 80% of sessions contain two and three check-in records, which indicates that a typical activity consists of more than one but less than four actions. For example, an activity of gathering friends may include the successive actions of eating, massage, and going to a bar. For our task of local recommendation, only the related and consistent check-in actions are useful. Therefore, we will focus on usage mining on the session level. By aggregating the inter check-in time within individual sessions, the estimated average inter check-in time $\bar{T}_{inter} = 38.9 min$.

*3.1.3. Statistic III: Inter Check-In Distance & #POI Candidate.* Analog to inter check-in time, inter check-in distance indicates the geographic distance between two successive check-in records from the same user. Figure 5 plots the distribution of inter check-in distance within individual sessions in Shanghai. It shows that the inter check-in distances span from 0.1km to 1.5km, with the average inter check-in distance $\bar{D}_{inter} = 0.33\ km$.

To estimate the number of POI candidates in exploring the next check-in destination, we calculate the POI density in Shanghai as $361/km^2$. Therefore, assuming that $\bar{D}_{inter} = 0.33\ km$ is the maximum acceptable distance for next action, we attempt to recommend the next POI within the range of the radius $\bar{D}_{inter}$ (centered at the user's current location), and the number of POI candidates will be 108 ($\pi \times 0.33^2 \times 361$). This provides a good understanding for the challenge of the local recommendation problem. However, since the interdistance between successive actions varies significantly with travel vehicles and moving speed, this is a very coarse estimation. In practice, users

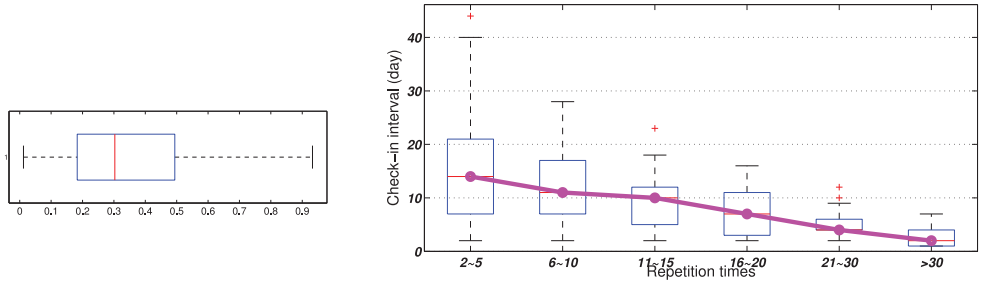


Fig. 6. Observation I: User repetition probability (left) and check-in interval versus repetition times distribution (right).

are willing to explore distant but interesting POIs for next destination, and it would be unreasonable to restrict the POI candidates within a fixed area. Therefore, the actual POI candidate number in practical local recommendation should be much larger.

### 3.2. Observations

*3.2.1. Observation I: POI Check-In Repeats a Lot.* Since people usually conduct activities around several fixed areas (e.g., home and company) and have personalized preferences over certain POIs, we are interested in the extent to which people check in to repetitive POIs. We examine the proportion of repetition in users' check-in records by calculating the ratio of unique POIs and define the repetition probability as

$$Repe(u) = 1 - \frac{\#uniquePOIs}{\#check-inrecords}. \tag{1}$$

The box plot on the left-hand side of Figure 6 shows the repetition probability distribution for the examined 2, 342 users. With the median value of repetition probability as 30.3%, we can see that a considerable proportion of POI user check-ins are repetitive.

The significant repetition frequency has long been discussed in music listening behaviors and exploited for music recommendation [Bradley 1971]. Note that most other types of recommender systems do not need to consider this observation for a solution (e.g., recommending books that have been bought or movies that have been watched makes little sense). However, the challenge in exploiting this observation is to repeat POIs at proper frequencies. To this end, we investigate the relationship between the POI repetition times and check-in time interval. The result is summarized on the right-hand side of Figure 6, where we can see that the check-in time interval and POI repetition times generally hold an inversely proportional relation. The median check-in interval is 14 days for those POIs that have been repeated 2 to 5 times, and when repetition times increase from 20 to 30, the median check-in interval drops to only 4 days. This indicates that when considering novelty in repetitive POI recommendation, the more times the POI has been checked-in, the sooner a user's novelty recovers, and thus with the higher probability it should be recommended.

*3.2.2. Observation II: Transition of POI Category Is Common.* To understand user mobile behavior patterns among the check-in records, the category of the check-in POI is very important, as the POI category carries more abstract activity information. In particular, we are interested to investigate how different POI categories succeed to one other in a session. We first examine how probable the latter check-in POI changes category with the former one (i.e., $s.C_k \neq s.C_{k+1}$). The category change probabilities for Beijing, Shanghai, and the overall collected dataset are shown on the left-hand side of Figure 7. We can see that all three curves ascend with the increase of inter check-in

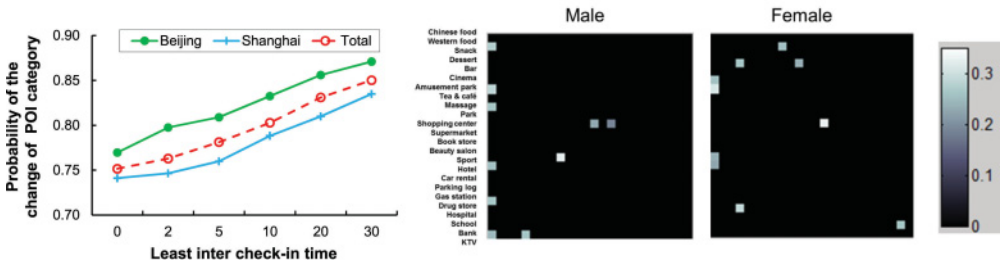


Fig. 7. Observations II & IV: POI category change probability (left) and popular transition patterns for male and female (right).

time. If we set the least inter check-in time to 10 minutes, the overall category change probability is around 0.8, indicating that four out of five check-in records have different POI categories with the former one.

We identify a significant transition pattern in check-in records at the POI category level, which is very different from traditional next-item recommendation problems. In Web page [Borges and Levene 2010], video [Davidson et al. 2010], music [Park et al. 2011], and product [Rendle et al. 2010]) recommendation problems, the recommended next items are expected to be similar to the previous one. For example, a user is much more likely to hear consequent jazz songs; however, he or she will be not likely to go to another restaurant shortly after visiting one. This observation motivates us to further investigate the POI category transition patterns and leads to the latter category-transition–based solution for local recommendation.

*3.2.3. Observation III: POI Category Transition Is Context Sensitive.* Typically, users will successively conduct a series of actions, such as visiting a snack store after shopping or going to a gas station or parking lot after obtaining a car rental. It is natural to assume that category transition within sessions follows certain patterns. In particular, we define the transition from category $i$ to $j$ as $s.C_k = i$ and $s.C_{k+1} = j$ $(1 \leqslant k \leqslant N_s - 1)$, where $s$ denotes the examined session, and $i$ and $j$ are the POI category of former and latter check-in records, respectively.

Since location provides key sensor context when performing local recommendation, we first investigate the POI category transition patterns in different locations. For simplicity, we examine the transition pattern at city level and compare between the category transition matrices for Beijing and Shanghai in the top portion of Figure 8. The gray scale on the right shows the corresponding probability values.[4] We can see that the transition patterns are much different for these two cities. For example, in Beijing, the transition pattern of *western food*→*bar* (element $(2, 5)$) is significant, whereas in Shanghai, the preferred POI category after *western food* is replaced by *tea & cafe* (element $(2, 8)$). This implies that the POI category transition is *location sensitive*.[5]

In addition to location, we further investigate the category transition patterns in different time slots. We split the time into 14 slots.[6] The category transition matrices for four selected time intervals are shown in the bottom portion of Figure 8. We can see that different time slots follow very different category transition patterns. For example,

[4]The dimension of the matrix is $24 \times 24$, where the value of element $(i, j)$ indicates the category transition probability from category $i$ to $j$.

[5]Actually, this observation is quite related to recent studies on culture and lifestyle investigation between different cities. In Yuan et al. [2013], the authors provide a very interesting discussion to this point.

[6]According to the check-in temporal distribution, we identified seven time slots in a day as 0:00-6:00, 6:00-8:00, 8:00-12:00, 12:00-13:00, 13:00-18:00, 18:00-20:00, and 20:00-24:00. We also treated the weekday and weekend separately so that the time is discretized into 14 time slots and coded by an integer index in [1, 14].

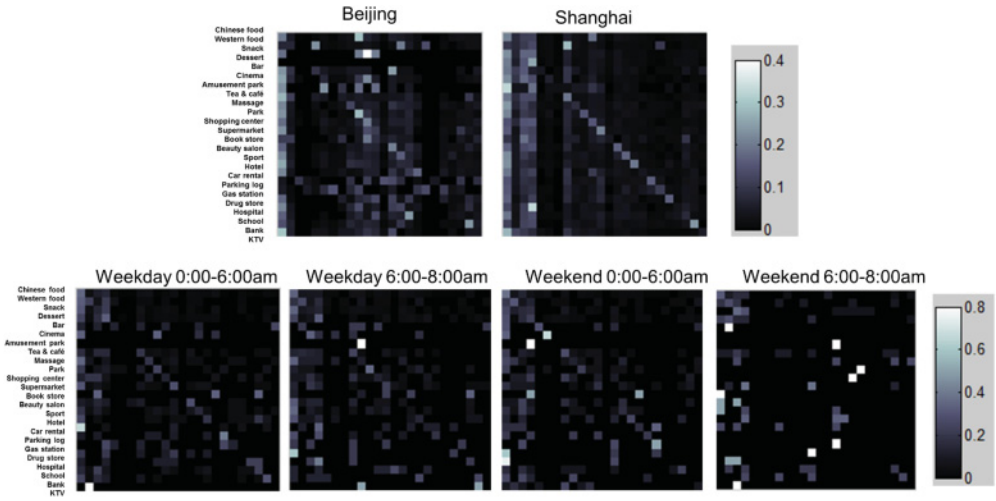


Fig. 8. Observation III: Category-transition matrices for different cities (top) and time slots (bottom).

time interval {*weekday 0:00-6:00*} has a distinctive transition pattern of *snack*→*hotel* (element (3, 16)), whereas in {*weekday 6:00-8:00*}, users tend to go to a *parking lot*, *bank*, or *school* after enjoying a grab-and-go *snack* (element (3, 18)(3, 23)(3, 22)). As another example, in {*weekday 0:00-6:00*}, a *hotel* is preferred after relaxing at a *bar* (element (5, 16)), whereas in {*weekend, 6:00-8:00*}, an overnight party is also a choice by going to *ktv* after a *bar* (element (5, 24)). Therefore, we can see that the POI category transition is also *time sensitive*.

*3.2.4. Observation IV: POI Category Transition Is Personalized.* After investigating the POI category transition pattern for different sensor contexts, it is interesting to ask whether it varies from user to user. We examine the 10 most popular category-transition patterns for female and male users and show the results on the right-hand side of Figure 7. Popular transition patterns indicate the typical sequential actions of users. The differences of popular transition patterns on the right-hand side of Figure 7 do not surprise us, as female and male users do conduct very different types of actions. For example, female users prefer to visit a *beautysalon* after a *shopping center* (element (11, 14)), whereas male users will be more likely to go to a *book store* or *gym* after a *shopping center* (element (11, 13)(11, 15)). This observation implies that category transition is *gender sensitive*—or to a bold claim—*personalized*.

The preceding statistics and observations from check-in usage mining inspire the design of a hierarchical, contextual, and sequential POI recommendation solution. In the next section, we propose the category-transition–based Markov POI ranking approach.

## 4. LOCAL RECOMMENDATION

We consider the local recommendation problem from a probabilistic ranking perspective. Specifically, the recommended POI sequences are ranked by $P(O_1, O_2, \ldots, O_N | O_0; U, L, T)$, where $U, L, T, O_0$ denote user, location, time, and POI of the current check-in action, and $O_1, O_2, \ldots, O_N$ is the recommended POI sequence including $N$ POIs. Inspired by the transition observation from check-in usage mining,

we formulate the target probability by POI transition probability:

$$\begin{aligned} &P(o_1, o_2, \ldots, o_N | o_0; u, l, t) \\ &\quad = P(o_1|o_0; u, l, t) \cdot R(o_1, u) \times P(o_2, \ldots, o_N | o_1, o_0; u, l, t) \cdot R(o_2, \ldots, o_N, u), \end{aligned} \tag{2}$$

where $P(o_j|o_i; u, l, t)$ is the transition probability from POI $o_i$ to $o_j$ under sensor context $l, t$ for user $u$, and $R(o, u)$ denotes the memory retention of user $u$ on POI $o$ (detailed in Section 4.3). In the following, we will interpret the formulation along three components of hierarchical, contextual, and sequential.

### 4.1. Hierarchical

The 1-order POI transition probability $P(o_j|o_i; u, l, t)$ is the core element in the formulation. Observations in Section 3 show significant category transition patterns in the successive check-in actions. Inspired by this, and to alleviate the challenge of huge POI candidates, we propose a category-transition–based hierarchical strategy to instantiate the POI transition probability as

$$\begin{aligned} P(o_j|o_i; u, l, t) &= (c(o_j)|o_i; u, l, t) \cdot P(o_j|c(o_j); u, l, t) \\ &\propto P(c(o_j)|c(o_i); u, l, t) \cdot P(o_j|c(o_j); u, l, t), \end{aligned} \tag{3}$$

where $c(o_i)$ is the category of POI $o_i$.

The probability of visiting next-POI $o_j$ from $o_i$ is first decomposed in two parts: (1) $P(c(o_j)|o_i; u, l, t)$—transition probability, the probability of conducting action denoted by $o_j$'s category $c(o_j)$ after the current check-in POI $o_i$, and (2) $P(o_j|c(o_j); u, l, t)$—POI popularity, the probability of choosing $o_j$ given its category $c(o_j)$. For example, imagine that Emily is now checking in to *Walmart* ($o_i$), if we know that she likes to visit a *cafe* ($c(o_j)$) after shopping at a *supermarket* ($c(o_i)$), and she prefers *Starbucks* ($o_j$) of all cafe stores, then we are confident to recommend a *Starbucks* nearby as her next-POI choice. Note that after the decomposition, we make approximation in the second step by assuming that the transition probability $P(c(o_j)|o_i; u, l, t)$ is mainly decided by the transition patterns between POI categories. This is reasonable in that it is the current action of shopping at a supermarket that decides the plan of going to a cafe—whether it is *Walmart* or *Carrefour* makes little difference. After approximation, the resultant $P(c(o_j)|c(o_i); u, l, t)$ is actually category transition probability, with better interpretation and easy calculation.

Before continuing to introduce how to calculate $P(c(o_j)|c(o_i); u, l, t)$ and $P(o_j|c(o_j); u, l, t)$, we first discuss that the hierarchical decomposition strategy also enables an efficient hierarchical organization for recommendation. Besides directly recommending POIs, we are interested to first recommend POI categories, which will help to narrow the POI candidates and facilitate users with an interactive coarse-to-fine experience. POI category recommendation is relatively easy, as it can be viewed as a subproblem of POI recommendation under our formulation. Similar to Equation (2), the probability of recommending a sequence of POI categories can be calculated as

$$P(c_1, c_2, \ldots, c_N | o_0; u, l, t) = P(c_1|o_0; u, l, t) \times P(c_2, \ldots, c_N | c_1, o_0; u, l, t), \tag{4}$$

where $c_k$ is the recommended POI category in the $k$-th step. By integrating the POI category recommendation results, in practice we address the local recommendation problem into a two-level organization (see Figure 1). In the first level, the candidate POI categories are ranked based on Equation (4). When users select a specific category sequence $< c_i, c_j >$ from the recommended category sequences, in the second level the POI sequences with the categories $c_i, c_j$ are filtered and ranked for recommendation based on Equation (2).

### 4.2. Contextual

We have considered user and sensor contexts in Equation (2) to represent the short-term intent. However, due to the limited check-in records, it is not practical to explicitly calculate the probabilities under these contexts. We address this by introducing a CF solution. The idea of user collaboration is widely used for personalized recommendation [Adomavicius and Tuzhilin 2005; Zhuang et al. 2011; Sang et al. 2012]. In the context of local recommendation, since sensor contexts like time and location are continuous and discretization will inevitably lead to the problem of context granularity [Bolchini et al. 2007], we conduct CF on both user and sensor contexts.

Assuming that the generation of users and contexts are independent, we have the following formulation for the category transition probability $P(c(o_j)|c(o_i); u, l, t)$:

$$\begin{aligned} &P(c(o_j)|c(o_i); u, l, t) \\ &\quad = \sum_{u_k;l_m;t_n} P(c(o_j)|c(o_i); u_k, l_m, t_n) P(u_k|u) P(l_m|l) P(t_n|t) \\ &\quad \propto \sum_{u_k;l_m;t_n} P(c(o_j)|c(o_i); u_k, l_m, t_n) a(u, u_k) a(l, l_m) a(t, t_n), \end{aligned}$$

where $a(\cdot, \cdot)$ is a measurement function for the similarity between users or contexts. The intuition for the approximation is that the user, location, and time transition probability is proportional to their similarities.

The calculation of POI popularity $P(o_j|c(o_j); u, l, t)$ is similar. The intuition is that if users similar to $u$ under similar context to $l, t$ have checked in to a POI of a certain category, then $u$ also has the possibility to choose the POI in this category. Therefore, we have

$$\begin{aligned} &P(o_j|c(o_j); u, l, t) \\ &\quad \propto \sum_{u_k;l_m;t_n} P(o_j|c(o_j); u_k, l_m, t_n) a(u, u_k) a(l, l_m) a(t, t_n). \end{aligned}$$

The conditional probability $P(c(o_j)|c(o_i); u_k, l_m, t_n)$ and $P(o_j|c(o_j); u_k, l_m, t_n)$ can be computed by directly counting frequencies in the collected check-in dataset. In the rest of this section, we focus on modeling the similarity functions $a(\cdot, \cdot)$ for collaborative user and context recommendation. Two correlations are considered for similarity functions: prior correlation and history correlation. Using location similarity as an example, we have

$$a(l_i, l_j) = \lambda a_{prior}(l_i, l_j) + (1 - \lambda) a_{history}(l_i, l_j), \tag{5}$$

where $\lambda$ is the weighting parameters,[7] $a_{prior}$ is the prior similarity function, and $a_{history}$ is the history similarity function.

History correlation is easy to understand in that the similarities between locations can be estimated by examining the history check-in records within the locations. Specifically, we represent each location by the checked POIs. Each location representation is then converted into a fixed-length vector according to term frequency–inverse document frequency (*tf-idf* [Sebastiani 2002]). We define $tf(o_k, l)$ as the frequency that a POI $O$ checked in location $l$. Assuming that $df(o_k)$ indicates the overall frequency of a POI $o_k$ across all locations, we construct the feature vector of $l$ as

$$[\mathbf{x}]_k = tf(o_k, l) \frac{1}{df(o_k)}.$$

[7]We empirically set $\lambda = 0.5$, which equals the contribution of prior and history information.

Then we can calculate the history correlation between two locations $l_i$ and $l_j$ as

$$a_{history}(l_i, l_j) = \mathbf{x}_i \mathbf{x}_j / \sqrt{(\mathbf{x}_i \mathbf{x}_i{}^\top)(\mathbf{x}_j \mathbf{x}_j{}^\top)}. \tag{6}$$

Prior correlation is used to exploit the available prior information of users and contexts for improved similarity estimation. In particular, we consider user prior information of gender and residence, time prior information of the temporal partition for different time intervals, and location prior of its geometrical relation. The basic premise is that two users have high prior similarity if they are of the same gender or from the same residence, and two time intervals have high prior similarity if they are adjacent within days or both from weekends or weekdays. Note that an alternative complex measurement of prior information can also be explored [Sang and Xu 2012].

Leveraging the available profile information, we calculate the user prior correlation as

$$a^U_{prior}(u_i, u_j) := \begin{cases} 1 & \text{if } u_i = u_j \\ 0.5 & \text{if } u_i.G = u_j.G \text{ and } u_i.Re = u_j.Re \\ 0.2 & \text{if } u_i.G = u_j.G \text{ or } u_i.Re = u_j.Re \\ 0 & \text{otherwise} \end{cases}, \tag{7}$$

where $u.G$, $u.Re$ denote the gender and residence of user $u$.

The time prior correlation is defined based on the temporal relations:

$$a^T_{prior}(t_i, t_j) := \begin{cases} 1 & \text{if } t_i = t_j \\ 0.5 & \text{if } t_i - t_j = 1 \text{ or } t_i - t_j = -1. \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

The prior similarity between two locations is inversely proportional to their distance:

$$a_{prior}(l_i, l_j) := \sigma\sqrt{2\pi} \cdot f\big(d(l_i, l_j); 0, \sigma^2\big), \tag{9}$$

where $d(l_i, l_j)$ is the distance between the center of $l_i$ and $l_j$[8], $f(x; 0, \sigma^2)$ is Gaussian distribution with mean 0 and variance $\sigma$, and the multiplier $\sigma\sqrt{2\pi}$ is a normalization term to make the function value equals 1 when $x = 0$.

### 4.3. Sequential

In the preceding two sections, we introduced how to calculate the first term in Equation (2). In this section, we introduce the sequential probability $P(o_2, \ldots, o_N | o_1, o_0; u, l, t)$ and memory retention $R(o, u)$.

$P(o_2, \ldots, o_N | o_1, o_0; u, l, t)$ denotes the probability of recommending POI sequence $o_2, \ldots, o_N$ if user $u$ has selected POI $o_1$ given the context of $(o_0; l, t)$. Assuming that the transition between POIs follows the Markov assumption, we calculate this conditional probability as

$$\prod_{k=1}^{N} P(o_{k+1} | o_k; u, l(o_k), t + k \cdot \bar{T}_{inter}), \tag{10}$$

where $l(o_k)$ is the location of POI $o_k$, and $\bar{T}_{inter}$ is the average inter check-in time interval estimated in Statistic II. Since we have no accurate sensor context information for the check-ins to predict, the next check-in location is replaced by the location of the candidate POI $o_k.L$, and the next check-in time is estimated by adding $\bar{T}_{inter}$ to the

[8]Location is discretized into areas, with each area positioned with the GPS of its center.

current timestamp. Note that the Markov assumption fails when recommending a long sequence of POIs. However, according to the session length distribution on the right-hand side of Figure 4, most user activities contain no more than three actions. Our experiments demonstrate the effectiveness of this formulation on recommending two-POI sequences. Moreover, in practice, a user's interaction can be incorporated into the recommendation by iteratively updating the candidate POI and context with the actual check-in actions.

The memory retention function $R(o, u)$ is to address the repetition issue discovered in Observation I. Many recommendation models considering repetition do not take time into consideration [Lathia et al. 2010]; as a result, items visited a year ago and just now have the same impact on the current recommendation. Inspired by Hu and Ogihara [2011], we assume that the freshness of a particular POI decays immediately after checking in and then gradually recovers. In this context, we define retention function as

$$R(o, u) = 1 - e^{\frac{-t \cdot rep(o,u)}{Z}}, \tag{11}$$

where $t$ is the time elapsed since user $u$ checked in to POI $o$ the last time, $rep(o, u)$ is the repetition times of POI $o$ for $u$, and $Z$ is set as the maximum value among all repeated POIs for user $u$. $rep(o, u)$ is actually a parameter controlling the freshness recovery speed, with faster recovery having a higher $rep(o, u)$. The premise here is that the repetition times indicates a user's preference over the POI, and thus frequently visited POIs deserve higher ranks. This has been validated on the right-hand side of Figure 6 from Observation I.

Since the sequential POI recommendation is decomposed into several steps of 1-order transition probability, the recommendation easily drifts to the latter POI if the former POI is wrong. In this case, the memory retention function also plays a role in alleviating recommendation drift by regularizing the final probability with a user's long-term preference encoded in $rep(o, u)$. Therefore, the memory retention function contributes toward the balance between POI freshness and user preference.

## 5. EXPERIMENTS

### 5.1. Settings

We conducted experiments on the collected check-in data of Shanghai, which include $105,463$ unique check-in records and $11,483$ check-in sessions from $1,114$ users. The number of POIs used in the evaluation is $25,503$, which are used as candidates for recommendation in the test set. We randomly selected 311 users and their check-in sessions from May 21 to June 11 in 2011 as test data. The test set contains $1,265$ check-in sessions, denoted by $\mathcal{S}_t$.

For location discretization and coding, we geometrically divide Shanghai into 400 local areas, with each area spanning a $5 \times 5$ km$^2$ square area. According to the data analysis in Shanghai, successive check-in seldom happens at areas 20km away from the current check-in. Therefore, the variance parameter $\sigma$ in Equation (9) is set as 8 so that location prior similarity is rather small ($a_{prior} \approx 0.05$) when the two location areas are 20km apart.

We compared the performance among the following recommendation schemes:

—*CF based (CF)*: Ranking candidates based on their co-occurrences with the current POI, which is implemented according to Horozov et al. [2006].
—*POI-transition based (PT)*: POI transition-based recommendation without considering user or context information, which can be viewed as 1-order implementation of Borges and Levene [2010].

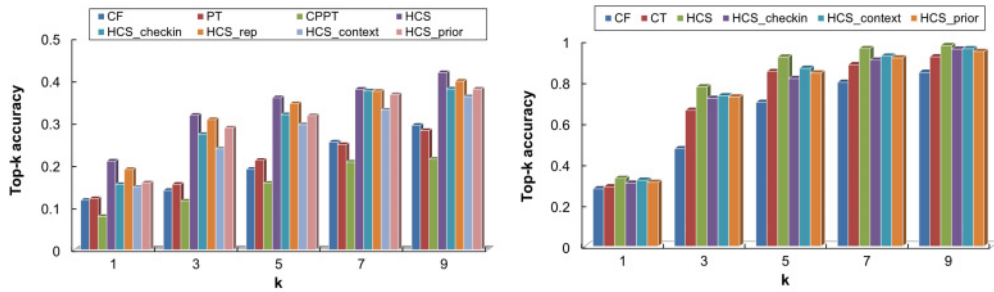


(a) POI recommendation (b) POI category recommendation

Fig. 9. The top-k accuracy for the examined methods.

—*Context and personalized POI transition based (CPPT)*: Contextual and personalized POI transition-centric recommendation, which directly computes the conditional probability $P(O_1|O_0; U, L, T)$.
—*Hierarchical, contextual, and sequential recommendation (HCS)*: The proposed method.
—*HCS_checkin*: The proposed method without considering current check-in record, which ranks based on $P(O_1, \ldots, O_N|U, L, T)$.
—*HCS_rep*: The proposed method without considering the memory retention function.
—*HCS_context*: The proposed method without considering context collaboration.
—*HCS_prior*: The proposed method without considering prior information in constructing the similarity functions.

### 5.2. Experiment I: Recommending POI and POI Category

For each test check-in session $s$, we use the location $s.L_0$ and time $s.T_0$ of its first check-in record as the context and employ the recommendation approach to recommend POI based on user identification $s.U$, current check-in POI $s.O_0$, and its POI category $s.C_0$. Top-k accuracy and mean absolute error (MAE [Witten and Frank 2005]) are used as the evaluation metrics. All of these 1, 265 test sessions ($N_s \geqslant 2$) constitute the test set $S_{t1}$ for recommending one POI. Denoting the candidate POI set as $\mathcal{O}_{t1} = \mathcal{O}$, $\psi(s.O_1)$ indicates the ranking position of $s.O_1$ in $\mathcal{O}_{t1}$ according to probabilistic POI ranking. Then, the top-k accuracy of POI recommendation is calculated by

$$Accuracy(\mathcal{S}_{t1}, \mathcal{O}_{t1}, k) = \frac{\sum_{s \in \mathcal{S}_{t1}} \mathbb{I}(\psi(s.O_1) \leqslant k)}{|\mathcal{S}_{t1}|}, \tag{12}$$

where $\mathbb{I}(\cdot)$ is indicator function returning 1 if it is true and 0 otherwise. We use MAE to access the average position of the ground-truth POI in the recommended list, which is

$$MAE(\mathcal{S}_{t1}, \mathcal{O}_{t1}) = \frac{1}{|\mathcal{S}_{t1}|} \sum_{s \in \mathcal{S}_{t1}} \psi(s.O_1) - 1. \tag{13}$$

To calculate the top-k accuracy and MAE for POI category recommendation, $O_1$ is replaced by $C_1$ and $\mathcal{O}_{t1}$ is replaced by $\mathcal{C}_{t1} = \mathcal{C}$.

Figure 9(a) shows the top-k accuracy for recommending one POI. Several observations can be drawn as follows. First, top-9 accuracy of two baseline methods, CF and PT, is only around 25%. CF and PT recommend POI based on the general inter-POI relationship and ignore the sensitivity to context and user, which results in inferior performance. This result demonstrates the need for contextual and personalized solutions. Second, CPPT has the worst performance among the evaluated methods. This

Table III. MAE of POI and POI Category Recommendation

| | CF | PT/CT | CPPT | HCS | HCS_checkin | HCS_rep | HCS_context | HCS_prior |
|---|---|---|---|---|---|---|---|---|
| POI Recommendation | 28.3 | 30.4 | 42.9 | **11.1** | 15.3 | 12.5 | 18.7 | 16.4 |
| Category Recommendation | 2.93 | 2.29 | – | **1.31** | 1.74 | – | 1.40 | 1.55 |

can be explained by two reasons: (1) due to the large number of POIs, the POI transition matrix is extreme sparse and it is impractical to explicitly calculate POI-level transitions; (2) different from other next-item recommendation problems, POI recommendation has close relation with action transfer, which is essentially a category-transition problem. This calls for the need for a hierarchical solution. Third, the fact that HCS significantly outperforms HCS_checkin validates the advantage of utilizing the current check-in action. The user's current check-in indicates the activity being conducted or to be conducted, which is important for understanding short-term user intent. Fourth, HCS_context produces lower accuracy than HCS. HCS_context considers collaboration only for the user factor and not for the context factor. It thus suffers from the sparsity issue and fails to make full use of the data. Fifth, another alternative, HCS_prior, also achieves inferior performance, as employing the prior information contributes to reducing the impact of context discretization. For example, check-in records at time interval of 12:00-13:00 should contribute to the recommendation of a check-in at 11:55.

We also evaluate the performance of POI category recommendation. Since category recommendation involves no POI transition, we remove the evaluated method of CPPT and replace PT with the category-transition–based recommendation method (CT). The results are shown in Figure 9(b). It is observed that the top-k accuracy of category recommendation follows similar patterns with that of POI recommendation. The proposed HCS consistently outperforms other baseline methods. The top-3 accuracy of HCS approaches 77%, which demonstrates its efficacy. MAE for different evaluation methods are shown in Table III. MAE serves as a rough estimation for the recommendation performance. For example, MAE of HCS for POI recommendation is 12.5, which means that the average position of ground-truth POI is 13.5 ($=12.5+1$). We can see that HCS produces the lowest MAE among the evaluated methods for both POI and POI category recommendation.

The experimental result shows the advantage of our proposed hierarchical, contextual solution for recommending one POI and POI category. This is important in two ways. On the one hand, it is difficult to predict the latter ones if the former recommended POI is not in accordance to the user's real choice. On the other hand, if the recommendation of the first POI is convincing and the user has a higher probability to follow it, in practical applications we can iteratively recommend the following POIs by incorporating user interaction into the loop.

### 5.3. Experiment II: Recommending POI Sequence

In the following, we evaluate the performance of recommending POI/POI category sequence. In 317 sessions out of the $1,265$ test sessions, there are no less than three check-in actions ($N_s \geqslant 3$), and they constitute the test set $S_{t2}$ for recommending two successive POIs. The candidate POI set is $\mathcal{O}_{t2} = \mathcal{O}^2$, and $\psi(s.O_1, s.O_2)$ indicates the position of $s.O_1, s.O_2$ in $\mathcal{O}_{t2}$. The top-k accuracy can thus be calculated as

$$Accuracy(\mathcal{S}_{t2}, \mathcal{O}_{t2}, k) = \frac{\sum_{s \in \mathcal{S}_{t2}} \mathbb{I}(\psi(s.O_1, s.O_2) \leqslant k)}{|\mathcal{S}_{t2}|}. \tag{14}$$

Table IV. MAE of POI Category Sequence Recommendation

| | CT | HCS | HCS_checkin | HCS_context | HCS_prior |
|---|---|---|---|---|---|
| Category Recommendation | 9.75 | **4.60** | 7.23 | 6.11 | 5.04 |

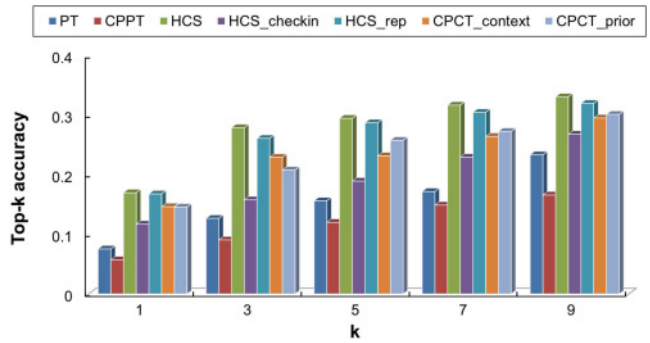


Fig. 10. The top-k accuracy for category-POI recommendation.

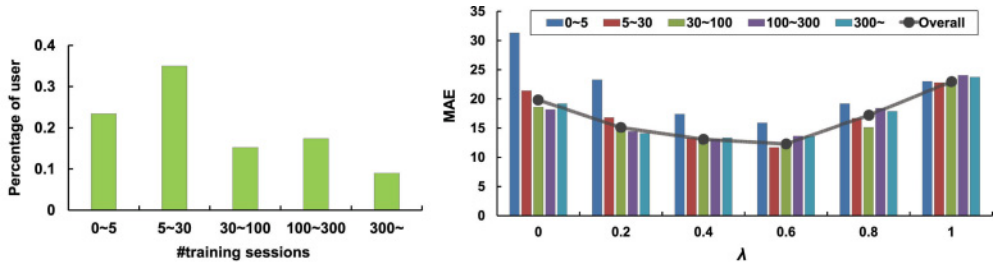


(a) User *#training sessions* distribution. (b) Sensitivity of MAE with different $\lambda$

Fig. 11. Experiment III: Sensitivity to prior.

MAE is

$$MAE(\mathcal{S}_{t2}, \mathcal{O}_{t2}) = \frac{1}{|\mathcal{S}_{t2}|} \sum_{s \in \mathcal{S}_{t2}} \psi(s.O_1, s.O_2) - 1. \tag{15}$$

Table IV shows MAE for recommending a 2-POI category sequence, where the lowest MAE = 4.6 is achieved by the proposed HCS. This indicates that the user has a high probability of finding the desired category sequence within the top six recommended results. To investigate the effectiveness of the two-level organization as shown in Figure 1, we examine the accuracy of recommending a POI sequence when the user chooses the desired category sequence $s.C_1, s.C_2$ in the first level. We denote the resultant top-k accuracy as category-POI accuracy, which is shown in Figure 10. Note that the number of candidate POI sequences $|\mathcal{O}|^2$ dramatically reduces to $N_{s.C_1} \times N_{s.C_2}$. The top-9 accuracy of HCS is around 0.3, which indicates that for one out of three trials, the top nine recommended POI sequences will include the user's actual choices for the following two actions.

### 5.4. Experiment III: Sensitivity to Prior

In Figure 11(a), we plot the distribution of #training sessions for each test user. About 20% of the test users have no more than five training sessions. We refer to the users with few history records as light users. On the one hand, a light user recommendation is a

widely recognized challenge [Agarwal and Chen 2009]. On the other hand, a light user recommendation is very important for the system viscosity. Therefore, in this section, we discuss the advantage of utilizing prior information for performance improvement on light user recommendation.

We integrate the prior information into user and context collaboration. The collaborative similarity is thus a combination of prior and history information. The 311 test users are separated into five groups according to their #training sessions. We evaluate the sensitivity of MAE of POI recommendation for different user groups to the prior weighting parameter $\lambda$.[9] The results are shown in Figure 11(b). It is seen that MAE of the lightest user group (0–5) is extremely high when $\lambda = 0$ and drops fast when $\lambda$ increases within [0, 0.6]. This demonstrates that user profile prior information is useful for light user recommendation. One interesting observation is that not only light users but also heavy users (user groups 100–300 and 300–) also benefit from the integration of user profile prior information.

### 5.5. Experiment IV: User Study

We developed a demo application of Activity Sensor on Windows Phone 7 devices, based on the proposed HCS local recommendation solution. The application is capable of recommending POI categories as well as POIs. We use the check-in records from both Beijing and Shanghai as the training dataset, and the candidate includes 32, 891 POIs with 24 POI categories. We evaluate the performance by a series of user studies.

We randomly identified 12 active users (6 from Beijing, 6 from Shanghai) who conducted more than 100 check-in records in the collected dataset. We then invited 12 participants to perform the user study in the form of role-playing: each participant was assigned a user ID and asked to think and act as the role of that user. In particular, the participants were provided the check-in history (each check-in record consists of "time," "location," "POI," and "POI category"), the check-in spatial distribution, temporal distribution, and POI category distribution to acquire an overall understanding of the user's check-in activity and preference. Moreover, the user's detailed profile information was also provided, including gender, residence, description, and even his or her homepage in Sina Weibo. Before the formal user study, the participants were encouraged to browse the Weibo homepage of the assigned users.

During the user study, each participant was asked to select four contexts (different time or location) from the users' check-in history. Based on the selected context and the given check-in record, four recommendation lists were presented to the participants:

—*POI_category*: POI category sequences ranked by the proposed HCS[10]
—*POI:HCS*: POI sequences ranked by HCS which belongs to the top ranked POI categories[11]
—*POI:PT*: POI sequences ranked by PT
—*POI:HCS_checkin*: POI sequences ranked by HCS_checkin.

Figure 12 shows an example of the four recommendation results given the same context and for the same user. Each participant was required to give a satisfaction score of each result on a scale of 1 to 5, in which 1 is worst and 5 is best. Therefore, there are $12 \times 4 \times 4 = 192$ scores in total. We computed the average score for each recommendation list across different participants. Table V shows the result. We can see that the score of POI category sequence recommendation is higher than POI

[9]HCS reduces to HCS_prior when $\lambda = 0$.
[10]Each category sequence consists of two successive POI categories.
[11]For POI sequence recommendation, we assume that the user clicks the rank-1 recommended category sequence.

(a) POI category sequence recommendation (b) POI sequence ranked by HCS (c) POI sequence ranked by PT (d) POI sequence ranked by HCS_checkin

Fig. 12. User study: The context is 2:30 pm on Friday, and the female user currently checks in to a beauty salon.

Table V. Subjective Evaluation Score from the User Study

| Scheme | POI Category | POI: HCS | POI: PT | POI: HCS_checkin |
|---|---|---|---|---|
| Score | 3.83 | 3.40 | 2.64 | 3.06 |

sequence recommendation, which indicates that the user's activity intent and preference are more easily captured at an abstract category level than specific POIs. Among the three POI sequence recommendation methods, the proposed HCS achieves better performance than HCS_checkin and PT, which is consistent with the objective evaluation results. This result further validates the effectiveness of the proposed approach for a better user experience. By combining the POI category and POI recommendation results, we are capable of providing users with an acceptable hierarchical, contextual, and sequential local recommendation application.

## 6. DISCUSSIONS

### 6.1. Data Are NOT Telling Everything

On the right-hand side of Figure 6, in Observation I we find that when the repetition times increases to more than 15, the check-in time interval drops fast, and when repetition times is more than 30, the median check-in time interval is only 2 days. Based on this data observation, we design the memory retention function in Equation (11) that freshness recovery of POI speeds up with the increase of #check-in times. By further analyzing this observation, for those POIs that have repetition times of more than 15, we find that a considerable percentage (21.5%) have a check-in interval at exactly 24 hours, which means that users check in to these POIs routinely every day at the same time.

In Figure 13, we plot the time and POI category distribution of these routinely everyday check-ins. The time distribution in 24 hours (left-hand side of Figure 13) shows two peaks around breakfast and lunch and has stable distribution between 5 and 10 pm. On the right-hand side of Figure 13, we find that the most frequent POI categories are *Chinese food*, *snack*, *dessert*, and *tea&cafe*. The two figures jointly indicate that users are likely to routinely check in to breakfast snack POIs and lunch/dinner restaurant POIs. We offer two discussions on this finding. First, more careful data analysis is needed for accurate usage mining. Ignoring the routine check-in cases results in the bias in estimating the relationship between check-in interval and repetition times, and thus leads to inappropriate solution design. Second, scenario awareness is key to practical applications. The goal of local recommendation is expected to make it easier

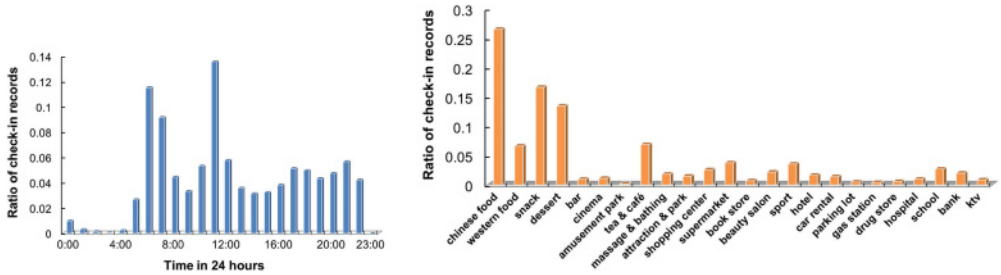


(a) Time distribution. (b) POI category distribution

Fig. 13. Routine everyday check-in statistics.

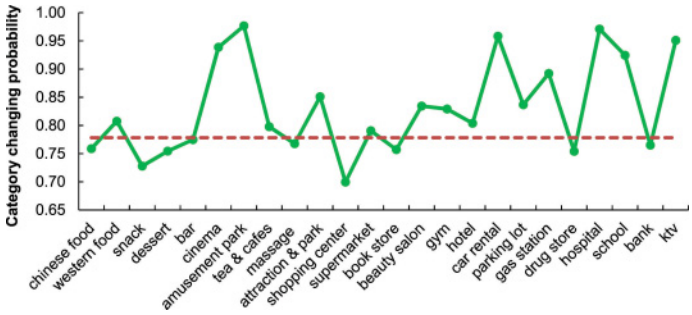


Fig. 14. Category change probability for each category.

for users to make decisions and help them explore novel POIs. The recommendation of routine POIs is of little use and should be excluded from performance evaluation.

### 6.2. Data Lies Sometimes

Although our proposed HCS recommendation solution has demonstrated its effectiveness in the evaluation, we find some phenomena hard to explain. For example, Figure 14 shows the category change probability for each category, where we can see that the category probabilities for *Chinese food* and *western food* are only around 75%. This indicates that about a quarter of users eat at successive Chinese or western restaurants within a short time, which is not reasonable.

We present two interesting discussions from this finding. First, we need to detect check-in spam. Where there is reward, there is spam. Due to the limited positioning accuracy and the large number of POIs, the current check-in service cannot differentiate spam from the standard check-in. Check-in spam detection is very important in future LBS services. In fact, the investigation on successive check-in records can motivate some solutions (e.g., it is impossible for one user to go to very distant POIs within a short time interval). Second, we need to separate *visit* from $pass-by$ check-in actions. Assuming that one user actually goes to a POI, there are still two types of check-in behaviors: visit and pass-by. Visit check-in indicates that the user pays a real visit to the POI (consumes or stays for a long time), whereas pass-by means that the user just passes through the POI without entering. It is easy to understand that only the visit check-in is meaningful. Observations in this article also offer some clues to the visit check-in identification; for example, within a short time, it is not common to eat at different restaurants, but it is quite possible to pass by different restaurants.

Besides, the visit check-in identification is category dependent. For example, visiting dozens of shopping centers one after the other is common for most shopaholics.

## 7. CONCLUSION

Location-based services are shifting from solving "simple" search tasks to addressing "advanced" recommendation tasks. However, the large number of available POIs and the complicated and changing user intent make local recommendation challenging. In this article, we conduct an in-depth usage mining on real-world check-in data and find inspiring observations to develop the hierarchical, contextual, and sequential solution for local recommendation. The successful implementation on a demo system validates its effectiveness. Future works include the following: (1) exploiting higher-order category-transition patterns by considering longer check-in sessions; (2) considering more attributes of POI (e.g., price range, ratings, comments) and social signals of users into recommendation; (3) integrating more context for recommendation, such as weather and user status (walking, driving, etc.); and (4) utilizing more advanced machine learning techniques to improve the current system [Zhou et al. 2013; Xu et al. 2014].

## REFERENCES

Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 6, 734–749.

Deepak Agarwal and Bee-Chung Chen. 2009. Regression-based latent factor models. In *Proceedings of KDD*. 19–28.

Ana O. Alves, Filipe Rodrigues, and Francisco C. Pereira. 2011. Tagging space from information extraction and popularity of points of interest. In *Proceedings of AmI*. 115–125.

Jie Bao, Yu Zheng, David Wilkie, and Mohamed F. Mokbel. 2013. A survey on recommendations in location-based social networks. Submitted to *Geoinformatica*.

Cristiana Bolchini, Carlo Curino, Elisa Quintarelli, Fabio A. Schreiber, and Letizia Tanca. 2007. A data-oriented survey of context models. *ACM SIGMOD Record* 36, 4, 19–26.

José Borges and Mark Levene. 2010. A comparison of scoring metrics for predicting the next navigation step with Markov model-based systems. *International Journal of Information Technology and Decision Making* 9, 4, 547–573.

Ian L. Bradley. 1971. Repetition as a factor in the development of musical preferences. *Journal of Research in Music Education* 19, 3, 295–298.

Jonathan Chang and Eric Sun. 2011. Location3: How users share and respond to location-based data on social. In *Proceedings of ICWSM*.

Kai-Ping Chang, Ling-Yin Wei, Mi-Yeh Yeh, and Wen-Chih Peng. 2011. Discovering personalized routes from trajectories. In *Proceedings of GIS-LBSN*. 33–40.

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 AAAI*.

Chen Cheng, Haiqin Yang, Michael R. Lyu, and Irwin King. 2013. Where you like to go next: Successive point-of-interest recommendation. In *Proceedings of IJCAI*. 2605–2611.

Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: User movement in location-based social networks. In *Proceedings of KDD*. 1082–1090.

Robert Cooley, Bamshad Mobasher, and Jaideep Srivastava. 1999. Data preparation for mining World Wide Web browsing patterns. *Knowledge and Information Systems* 1, 1, 5–32.

Justin Cranshaw, Eran Toch, Jason I. Hong, Aniket Kittur, and Norman M. Sadeh. 2010. Bridging the gap between physical location and online social networks. In *Proceedings of UbiComp*. 119–128.

James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, and Dasarathi Sampath. 2010. The YouTube video recommendation system. In *Proceedings of RecSys*. 293–296.

Huiji Gao, Jiliang Tang, and Huan Liu. 2012. gSCorr: Modeling geo-social correlations for new check-ins on location-based social networks. In *Proceedings of CIKM*. 1582–1586.

Marta C. Gonzalez, Cesar A. Hidalgo, and Albert-Laszlo Barabasi. 2008. Understanding individual human mobility patterns. *Nature* 453, 7196, 779–782.

Tzvetan Horozov, Nitya Narasimhan, and Venu Vasudevan. 2006. Using location for personalized POI recommendations in mobile environments. In *Proceedings of SAINT*. 124–129.

Yajie Hu and Mitsunori Ogihara. 2011. NextOne player: A music recommendation system based on user behavior. In *Proceedings of ISMIR*. 103–108.

Takeshi Kurashima, Tomoharu Iwata, Takahide Hoshide, Noriko Takaya, and Ko Fujimura. 2013. Geo topic model: Joint modeling of user's activity area and interests for location recommendation. In *Proceedings of WSDM*. 375–384.

Nicholas D. Lane, Dimitrios Lymberopoulos, Feng Zhao, and Andrew T. Campbell. 2010. Hapori: Context-based local search for mobile phones using community behavioral modeling and similarity. In *Proceedings of UbiComp*. 109–118.

Georgios Lappas. 2011. From Web mining to social multimedia mining. In *Proceedings of ASONAM*. 336–343.

Neal Lathia, Stephen Hailes, Licia Capra, and Xavier Amatriain. 2010. Temporal diversity in recommender systems. In *Proceedings of SIGIR*. 210–217.

Defu Lian and Xing Xie. 2011. Collaborative activity recognition via check-in history. In *Proceedings of LBSN*.

Bin Liu and Hui Xiong. 2013. Point-of-interest recommendation in location based social networks with topic and location awareness. In *Proceedings of SDM*. 396–404.

Anastasios Noulas, Salvatore Scellato, Cecilia Mascolo, and Massimiliano Pontil. 2011. An empirical study of geographic user activity patterns in foursquare. In *Proceedings of ICWSM*.

Sung Eun Park, Sangkeun Lee, and Sang Goo Lee. 2011. Session-based collaborative filtering for predicting the next song. In *Proceedings of CNSI*. 353–358.

Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of WWW*. 811–820.

Adam Sadilek, Henry A. Kautz, and Jeffrey P. Bigham. 2012. Finding your friends and following them to where you are. In *Proceedings of WSDM*. 723–732.

Jitao Sang, Tao Mei, Ying-Qing Xu, Chen Zhao, Changsheng Xu, and Shipeng Li. 2013. Interaction design for mobile visual search. *IEEE Transactions on Multimedia* 15, 7, 1665–1676.

Jitao Sang and Changsheng Xu. 2012. Right buddy makes the difference: An early exploration of social relation analysis in multimedia applications. In *Proceedings of MM*. 19–28.

Jitao Sang, Changsheng Xu, and Jing Liu. 2012. User-aware image tag refinement via ternary semantic analysis. *IEEE Transactions on Multimedia* 14, 3–2, 883–895.

Fabrizio Sebastiani. 2002. Machine learning in automated text categorization. *ACM Computing Surveys* 34, 1, 1–47.

Panagiotis Symeonidis, Antonis Krinis, and Yannis Manolopoulos. 2013. GeoSocialRec: Explaining recommendations in location-based social networks. In *Proceedings of ADBIS*. 84–97.

Xiaozhe Wang, Ajith Abraham, and Kate A. Smith. 2005. Intelligent Web traffic mining and analysis. *Journal of Network and Computer Applications* 28, 2, 147–165.

Ling-Yin Wei, Yu Zheng, and Wen-Chih Peng. 2012. Constructing popular routes from uncertain trajectories. In *Proceedings of KDD*. 195–203.

Ian H. Witten and Eibe Frank. 2005. *Data Mining: Practical Machine Learning Tools and Techniques* (2nd ed.). Morgan Kaufmann, San Francisco, CA.

Chang Xu, Dacheng Tao, and Chao Xu. 2014. Large-margin multi-view information bottleneck. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 8, 1559–1572.

Mao Ye, Peifeng Yin, and Wang-Chien Lee. 2010. Location recommendation for location-based social networks. In *Proceedings of GIS*. 458–461.

Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik Lun Lee. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *Proceedings of SIGIR*. 325–334.

Nicholas Jing Yuan, Fuzheng Zhang, Defu Lian, Kai Zheng, Siyu Yu, and Xing Xie. 2013. We know how you live: Exploring the spectrum of urban lifestyles. In *Proceedings of COSN*. 3–14.

Vincent Wenchen Zheng, Bin Cao, Yu Zheng, Xing Xie, and Qiang Yang. 2010a. Collaborative filtering meets mobile recommendation: A user-centered approach. In *Proceedings of AAAI*.

Vincent Wenchen Zheng, Yu Zheng, Xing Xie, and Qiang Yang. 2010b. Collaborative location and activity recommendations with GPS history data. In *Proceedings of WWW*. 1029–1038.

Yu Zheng and Xing Xie. 2011. Learning travel recommendations from user-generated GPS traces. *ACM Transactions on Intelligent Systems and Technology* 2, 1, Article No. 2.

Tianyi Zhou, Wei Bian, and Dacheng Tao. 2013. Divide-and-conquer anchoring for near-separable nonnegative matrix factorization and completion in high dimensions. In *Proceedings of ICDM*. 917–926.

Jinfeng Zhuang, Tao Mei, Steven C. H. Hoi, Ying-Qing Xu, and Shipeng Li. 2011. When recommendation meets mobile: Contextual and personalized recommendation on the go. In *Proceedings of UbiComp*. 153–162.

Received December 2013; revised September 2014; accepted September 2014

## QUERY

**Q1:** AU: Please provide complete mailing and email addresses for all authors.