Probabilistic Sequential POIs Recommendation via Check-In Data [*]

Jitao Sang [1,3], Tao Mei [2], Jian-Tao Sun [2], Changsheng Xu [1,3], Shipeng Li [2]
[1] National Lab of Pattern Recognition, Institute of Automation, CAS, Beijing 100190, P. R. China
[2] Microsoft Research Asia, Beijing 100080, P. R. China
[3] China-Singapore Institute of Digital Media, Singapore 119613
{jtsang, csxu}@nlpr.ia.ac.cn, {tmei, jtsun, spli}@microsoft.com

ABSTRACT

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. While existing research predominantly focuses on one-step recommendation—recommending the next single activity according to current context, this work moves one step beyond by recommending a series of activities, which is a package of sequential Points of Interest (POIs). The recommended POIs are not only *relevant* to user context (i.e., current location, time, and check-in), but also *personalized* to his/her check-in history. We presents a probabilistic approach, which is highly motivated from a large-scale commercial mobile check-in data analysis, to ranking a list of sequential POI categories (e.g., "Japanese food" and "bar") and POIs (e.g., "I love sushi"). The approach enables users to plan consecutive activities on the move. Specifically, the probabilistic recommendation approach estimates the transition probability from one POI to another, conditioned on current context and check-in history in a Markov chain. To alleviate the discritization error and sparsity problem, we further introduce context collaboration and integrate prior information. Experiments on over $100k$ real-world check-in records and $20k$ POIs validate the effectiveness of the proposed approach.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*

General Terms

Algorithms, Experimentation, Performance.

Keywords

Location-based services, sequential POIs ranking, check-in record.

[*]This work was performed at Microsoft Research Aisa.

ACM SIGSPATIAL GIS '12, Nov. 6-9, 2012. Redondo Beach, CA, USA

Figure 1: The interface of sequential POI recommendation application developed based on the proposed approach. A user named "Emily" checks in a shopping mall at the context of (a), and gets recommendation of (b) sequential POI categories (e.g., cafe→restaurant) and (c) the corresponding packages of sequential POIs when selecting a specific POI category sequence in (b). The recommendation of sequential POI categories and POIs rather a single POI is a metaphor of "plan activities in real life" and thus more natural to mobile users. Users will have much less interactions with the phone to complete their tasks.

1. INTRODUCTION

The ubiquitous mobile devices have revolutionized the way people surf information and make decisions. People are using their mobile phones to accomplish tasks anytime and anywhere. It is therefore important to understand user intent and provide context-aware and personalized recommendation services on the go. This is challenging because user intent could be "complicated"—in many cases, a user does not have a very explicit intent in mind, but just wants someone to suggest or plan activities for him/her. Existing research has predominantly focused on recommendation by understanding the "simple" intent. For example, previous research attempts to predict user's search intent by recognizing meaningful entities from a textual or voice query [3], or recommend a POI according to user's context (location and time) and community behaviors [6]. However, understanding the "complicated" user intent mentioned above still remains challenging.

We are investigating in this work if your mobile phone can facilitate planning your activities, by conducting contextual and personalized recommendation of sequential POIs. This is, to the best of our knowledge, very distinctive to most existing mobile recommendation applications in that we are trying to provide a package of several sequential activities (i.e., POIs) conditioned on current context, rather than an individual POI. Note that *sequential* POI recommendation is reasonable and practical. On one hand, users are likely to conduct consecutive actions, e.g., having dinner first then going to Karaok, and sequential POI recommendation accords with their behavior patterns. On the other hand, the required sequential POI correlation can be inferred from user activity data, which is statistically convinced. Moreover, we validate in this paper that users' current action indicates the activity he/she is conducting or will conduct, which is important for complicated user intent understanding.

We formulate the problem of planning activities as given the context (such as time and location) and user's current check-in action [1], we recommend a list of sequential POIs to be visited (as shown in Figure 1). Typical activity planning consists of a series of actions, where each action can corresponds to one single check-in POI. Therefore, we intend to facilitate mobile users by recommending a series of succeeding POIs. For example, the recommendation of "a tea break in some nearby cafe" and then "a decent Western restaurant" to a lady who checks in a shopping mall would be a nice-to-have 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.

The contributions of this paper can be summarized as follows:

- We address the complicated user intent understanding problem by moving one step beyond most existing recommendation approaches—recommending a series of sequential POI categories and POIs, rather than an individual category or POI. Users' current action is explicitly considered and modeled.
- We design a probabilistic approach by incorporating context collaboration and prior information in a Markov chain, which offers a simple yet effective solution to contextual and personalized recommendation.

2. APPROACH

We collected a check-in dataset from Dianping [2], the check-in website cooperating with Sina Weibo. The collected check-in records are located at Beijing and Shanghai, from Jan 7 to June 11 in 2011. The overall number of check-in records is $152,154$, issued by $2,342$ unique users. For the duplicate records where two consecutive check-in POIs are identical, we removed the latter. After filtering, we obtained $128,937$ records in total, with 55 average records per user. $32,891$ unique POIs were identified from these records, resulting in 24 POI categories with public facilities, entertainments, life services, and so on.

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 record are listed in Table 1. We represent each check-in record as a 5-dimensional tuple:

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

where $U \in \mathcal{U}$ is the user ID [3], L and T are the context (i.e., location and time when R was generated), $O \in \mathcal{O}$ is the check-in POI, and C is the category of O. Besides the check-in records, we also collected the gender and residence information of each user, as well as the attributes of each POI, such as the name, category, GPS location, and address. In this work, we use "session" to represent an activity with a series of consistent check-in actions. Based on the above discussions, each session can be represented as:

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

Table 1: Example of check-in record.

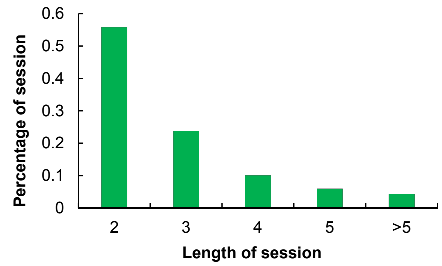

Figure 2: The distribution of session length.

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

where N_s is the length of session, 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. As a result, $13,078$ sessions are extracted. The distribution of session lengths is shown in Figure 2. We can observe that around 80% sessions containing two and three check-in records, which demonstrates that a typical "activity planning" consists of more than one but less than four activities. By aggregating the inter check-in time within individual sessions, the estimated average inter check-in time $\bar{T}_{inter}$ is 38.9 min.

We consider the recommendation problem from a ranking perspective. Specifically, the recommended POI sequences are ranked by $P(O_2, O_3, \cdots, O_N | O_1; U, L, T)$, where U, L, T, O_1 denote user, location, time, and POI of the current check-in action, $O_2, \cdots, O_N$ is the recommended POI sequence including $(N-1)$ POIs. Three reasons are considered to simplify the conditional probability based on Markov assumption: (1) according to the session length distribution in Figure 2, most user activities contain no more than three actions; (2) in practical applications, user's interaction can be incorporated into the recommendation by iteratively updating the candidate POI and context with user's real check-in actions. Following the Markov assumption, the target conditional probability is formulated as:

$$\begin{aligned} &P(o_2, \cdots, o_N | o_1; u, l, t) = P(o_2 | o_1; u, l, t) \times \\ &\prod_{k=2}^{N-1} P\Big(o_{k+1} | o_k; u, l(o_k), t + (k-1)\bar{T}_{inter}\Big). \end{aligned} \tag{1}$$

We propose a category transition-centric formulation [4] and assume the probability of visiting next-POI o_j is determined by two components: 1) 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) POI popularity—the probability of choosing o_j given its category $c(o_j)$. Therefore, the target conditional probability can be decomposed as:

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

[1] In social networking services, user check in to a physical place (i.e., POI) and share their locations with friends. User check-in action is a good indicator of the activity being conducted and naturally bridges user intent and behaviors in the physical world.

[2] http://www.dianping.com/

[3] The user information is fully anonymous.

[4] Our data analysis on check-in logs shows significant transition patterns at the category level.

where $P\big(c(o_j)|c(o_i);u,l,t\big)$ denotes category transition probability conditioned on (u,l,t), and $P\big(o_j|c(o_j);u,l,t\big)$ denotes POI popularity.

Due to the limited check-in history, it is not practical to explicitly calculate the category transition probability $P\big(c(o_j)|c(o_i);u,l,t\big)$. We introduce user and context collaborative recommendation by considering the correlations between contexts as well as users to overcome the sparsity and discretization problems:

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

where $a(\cdot,\ \cdot)$ is a similarity measurement function between users or contexts. The intuition for the approximation is that the user, location and time transition probability is proportional to their similarities. To alleviate the influence of inter check-in context, we also introduce context and user collaboration in calculating the POI popularity:

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

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.

Two correlations are considered on modeling the similarity functions $a(\cdot,\cdot)$: prior correlation and history correlation. History correlation is easy to understand that the users' or contexts' similarity can be represented by the history check-in records of the users or within the contexts. We use the prior correlation to model the prior information for users and contexts, where user prior information includes their gender and residence, time prior information indicates the temporal partition for different time intervals and location prior means the geometrical information. Using location similarity as example, we have

$$a(l_i,l_j)=\lambda a_{prior}(l_i,l_j)+(1-\lambda_r)a_{history}(l_i,l_j) \qquad (2)$$

where λ is the weighting parameters, a_{prior} is the prior similarity function, and $a_{history}$ is the history similarity function.

We utilize heuristic methods for the prior similarity calculation: the prior similarity between two locations is inversely proportional to their distance, 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 day or both from weekend or weekday. Since transition probability and POI popularity focus on category transition and POI probability inside category, we construct different functions for the computation of their history similarities. Specifically, for transition probability history similarity, we represent each user or context by category-transition history for transition probability similarity, and check-in record history for POI popularity similarity. Each user or context is then converted into a fixed-length vector according to term frequency-inverse document frequency (*tf-idf* [4]). Here "document" indicates user U, location area L or time interval T, while "term" refers to either a category transition pattern $< C_i, C_j >$ or a POI O.

3. EXPERIMENTS

We conducted experiments on the collected check-in dataset. 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 each test check-in session s, we use the location $s.L_1$ and time $s.T_1$ of its first check-in record as the context and employ the recommendation approach to recommend POIs based on user identification $s.U$, current check-in POI $s.O_1$ and POI category $s.C_1$. We utilize top-k accuracy and mean absolute error (MAE [5]) as the evaluation metrics. We consider evaluating the performance on situations of recommending one and two POIs. All these $1,265$ test sessions ($N_s \geqslant 2$) constitute the test set S_{t1} for recommending one POI. 317 sessions out of the $1,265$ test sessions contain no less than three check-in actions ($N_s \geqslant 3$), which constitute the test set S_{t2} for recommending two POIs. Denoting the candidate POI set for recommending two POIs as $\mathcal{O}_{t2}=\mathcal{O}^2$, $\psi(s.O_2,s.O_3)$ indicates the position of $s.O_2, s.O_3$ in $\mathcal{O}_{t2}$ according to probabilistic POIs ranking. Then, the top-k accuracy of recommending two POIs is computed by

$$Accuracy(\mathcal{S}_{t2},\mathcal{O}_{t2},k)=\frac{\sum_{s\in\mathcal{S}_{t2}}\mathbb{I}(\psi(s.O_2,s.O_3)\leqslant k)}{|\mathcal{S}_{t2}|} \qquad (3)$$

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

$$MAE(\mathcal{S}_{t2},\mathcal{O}_{t2})=\frac{1}{|\mathcal{S}_{t2}|}\sum_{s\in\mathcal{S}_{t2}}\psi(s.O_2,s.O_3)-1 \qquad (4)$$

To calculate the top-k accuracy and MAE for one POI recommendation, $\mathcal{S}_{t2}$ is replaced by $\mathcal{S}_{t1}$ and $\mathcal{O}_{t2}$ is replaced by $\mathcal{O}_{t1}=\mathcal{O}$.

We compared the performance among the following recommendation schemes:

- Co-occurrence based (**CO**): recommending candidates that most co-occur with current POI, which is implemented as [2].
- POI-transition based (**PT**): POI transition centric recommendation without considering user or context information, which can be regarded as an one-order implementation of [1].
- Context and personalized POI transition based (**CPPT**): contextual and personalized POI transition centric recommendation, computing the conditional probability $P(O_2|O_1;U,L,T)$.
- Context and personalized POI category transition based (**CPCT**): POI category transition centric recommendation, which is our proposed approach in this work.
- **CPCT_context**: POI category transition centric recommendation without considering context collaboration.
- **CPCT_prior**: POI category transition centric recommendation without considering prior information in constructing similarity functions.
- **CPCT_checkin**: POI category transition centric recommendation without considering current check-in record, which ranks based on $P(O_2,\cdots,O_N|U,L,T)$.

Figure 3(a) shows the top-k accuracy for recommending one POI. Several observations can be drawn as follows. (1) Top-9 accuracy of two baseline methods, CO and PT is only around 25%. CO and PT recommend POI based on general inter-POI relationship and ignore the sensitivity to context and user, which results in inferior performances. This result demonstrates the need for contextual and personalized methods. (2) CPCT_context produces lower accuracy than CPCT. CPCT_context only considers collaboration for the user factor and does not leverage context collaboration, which subjects to the sparsity issue and fails to make full use of the data. (3) Another alternative, CPCT_prior, also achieves inferior performance than CPCT. Employing the prior information can reduce the impact of context discretization. For example, check-in records at

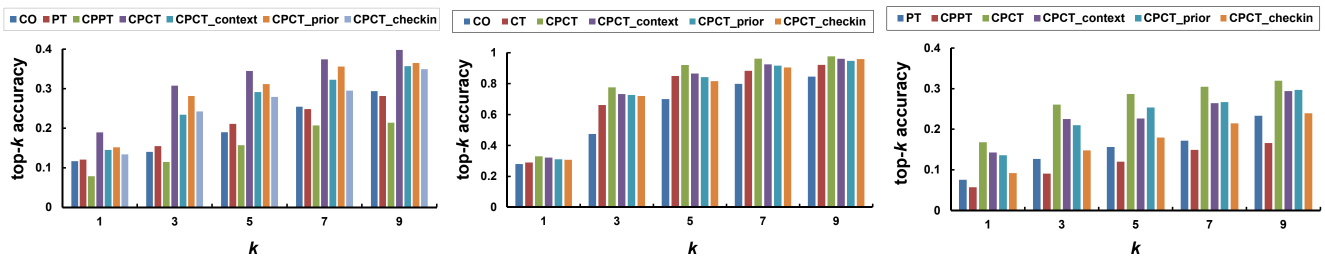

(a) one POI recommendation (b) one POI category recommendation (c) two category-POI recommendation

Figure 3: The top-k accuracy for the examined methods.

		CO	PT/CT	CPPT	CPCT	CPCT_context	CPCT_prior	CPCT_checkin
POI recommendation	one	28.3	30.4	42.9	**12.5**	19.1	17.8	20.7
Category recommendation	one	2.93	2.29	–	**1.31**	1.40	1.55	1.74
	two	–	9.75	–	**4.60**	6.11	5.04	7.23

Table 2: MAE of next-POI and next-POI category recommendation.

time interval of 12:00-13:00 should contribute to the recommendation to a check-in at 11:55. (4) The fact that CPCT significantly outperforms CPCT_checkin validates the advantage of utilizing the current check-in record. User's current action indicates the activity he/she is conducting or will conduct, which is important for user intent understanding and activity plan.

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 category-transition centric recommendation method (CT). The results are shown in Figure 3(b). It is observed that the top-k accuracy of POI category recommendation follows similar patterns with that of POI recommendation. The proposed method consistently outperforms the baseline methods. The top-3 accuracy of CPCT achieves 77%, which demonstrates its efficacy. The MAE are shown in Table 2. MAE serves as a rough estimation for the recommendation performance. For example, the MAE of CPCT for recommending one POI is 12.5, which means the average position of ground-truth POI is 13.5 ($= 12.5 + 1$). We can see CPCT produces the lowest MAE among the evaluated methods for both POI and POI category recommendation.

As shown in Table 2, the lowest MAE for recommending two POI categories is 4.6, which is achieved by the proposed CPCT. This demonstrates that user has high probability to find desired category sequence within the top 6 recommended results. Along the two-level organization as shown in Figure 1, we are interested to the accuracy of recommending POI sequence when user chooses the desired category sequence $s.C_2, s.C_3$ in the first level. We denote its top-k accuracy as category-POI accuracy, which is shown in Figure 3(c). Note the number of candidate POI sequences $|\mathcal{O}|^2$ dramatically reduces to $N_{s.C_2} \times N_{s.C_3}$. The top-9 of CPCT is around 30%, which means for one out of three trials, the nine top recommended POI sequences may include the users' actual choices for the following two actions. To summarize, the experimental results validate the effectiveness of the proposed approach for recommending one and two POIs and leaves room for improvement in the future work.

4. CONCLUSION

In this paper, we explore the potential of location-based service to cope with an advanced recommendation problem—activity plan, which is to suggest a package of sequential activities related to user context and interest. First, it is shown that check-in record is effective in user modeling and activity recommendation. Second, we found distinct category transition phenomenon in consecutive actions from a commercial check-in dataset, which has proved its significance in recommending sequential activities.

The future works include: 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, etc.) into recommendation, 3) leveraging social signals for better understanding user preference and thus improving recommendation performance, and 4) integrating more context for recommendation, e.g., weather, user status (walking, driving, etc.).

5. ACKNOWLEDGEMENT

This work was supported in part by National Program on Key Basic Research Project (973 Program, Project No. 2012CB316304) and the National Natural Science Foundation of China (Grant No. 90920303, 61003161).

6. REFERENCES

[1] J. Borges and M. Levene. 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, 2010.

[2] J. Davidson, B. Liebald, J. Liu, P. Nandy, T. V. Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, B. Livingston, and D. Sampath. The youtube video recommendation system. In *Proc. RecSys*, pages 293–296, 2010.

[3] N. D. Lane, D. Lymberopoulos, F. Zhao, and A. T. Campbell. Hapori: context-based local search for mobile phones using community behavioral modeling and similarity. In *Proc. UbiComp*, pages 109–118, 2010.

[4] F. Sebastiani. Machine learning in automated text categorization. *ACM Comput. Surv.*, 34(1):1–47, 2002.

[5] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann Publishers, San Francisco, CA, 2nd edition, 2005.

[6] J. Zhuang, T. Mei, S. C. H. Hoi, Y.-Q. Xu, and S. Li. When recommendation meets mobile: contextual and personalized recommendation on the go. In *Proc. Ubicomp*, pages 153–162, 2011.