

Personalized Point-of-Interest Recommendation by Mining Users' Preference Transition

Xin Liu
École Polytechnique Fédérale
de Lausanne
Lausanne, Switzerland
x.liu@epfl.ch

Yong Liu
Nanyang Technological
University
Singapore
liuy0054@e.ntu.edu.sg

Karl Aberer
École Polytechnique Fédérale
de Lausanne
Lausanne, Switzerland
karl.aberer@epfl.ch

Chunyan Miao
Nanyang Technological
University
Singapore
ascymiao@ntu.edu.sg

ABSTRACT

Location-based social networks (LBSNs) offer researchers rich data to study people's online activities and mobility patterns. One important application of such studies is to provide personalized point-of-interest (POI) recommendations to enhance user experience in LBSNs. Previous solutions directly predict users' preference on locations but fail to provide insights about users' preference transitions among locations. In this work, we propose a novel category-aware POI recommendation model, which exploits the transition patterns of users' preference over location categories to improve location recommendation accuracy. Our approach consists of two stages: (1) preference transition (over location categories) prediction, and (2) category-aware POI recommendation. Matrix factorization is employed to predict a user's preference transitions over categories and then her preference on locations in the corresponding categories. Real data based experiments demonstrate that our approach outperforms the state-of-the-art POI recommendation models by at least 39.75% in terms of recall.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Information filtering

General Terms

Algorithms, Performance

Keywords

Location-Based Social Networks, Point-of-Interest, Recommendation, User Preference

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM'13, Oct. 27–Nov. 1, 2013, San Francisco, CA, USA.
Copyright 2013 ACM 978-1-4503-2263-8/13/10 ...\$15.00.

1. INTRODUCTION

By aggregating the check-in data of millions of users, the LBSN has become a promising platform for studying users' online activities and mobility patterns. A lot of applications can be built on the basis of such studies (e.g., monitoring disease spread and traffic congestion), among which recommending personalized POIs to individual users is of particular interest, and has recently received a lot of research attention. For example, Ye et al. [10] and Cheng et al. [1] employed collaborative filtering to recommend POIs, taking into account geographical and social influence.

However, most existing approaches directly predict a user's preference on a location but fail to deeply investigate the essence of transitions of a user's preference on locations (e.g., why a user checks-in at a restaurant after work on Monday but checks-in at a stadium on Friday? Do such preference transitions have any predictable patterns? What factors influence a preference transition?). Failing to handle this important issue leads to inaccurate POI recommendation, since it is very challenging to accurately predict the next location from a huge candidate pool, and in particular, when input information is sparse, which is very common in LBSNs.

In order to improve the accuracy of POI recommendation, we investigate the patterns of users' preference transitions over location categories. In LBSNs, POIs are typically organized by a hierarchical category tree. For instance, Foursquare offers a 3-level category hierarchy where top level consists of *night life*, *food*, etc., and lower level consists of *bar*, *pub*, *Japanese food*, *coffee*, etc. Intuitively, such category information indicates the semantic meaning of a location, so a user's check-in at a given location reflects her preference on the corresponding category. Following previous work [4], we assume users' next check-in locations depend on their current locations. By studying patterns of users' preference transitions over location categories, we are able to infer the categories of their next check-in locations. And since bounded to specific categories, POI recommendations can be more accurate compared to traditional solutions where no category information is considered.

Theoretically, if a user has sufficient historical check-in data, her preference transitions can be reliably predicted [2, 7]. However, in a real world LBSN, users' check-in data is of-

ten sparse. Furthermore, even if a user has enough check-in data, she often presented new preference transitions. Such issues make traditional solutions (e.g., Markov model) ineffective for preference transition prediction. Therefore, we apply collaborative filtering to learn a user’s transition patterns by leveraging other *similar* users’ transition patterns. Similar users are clustered based on their check-in frequency in different categories, taking into account temporal effects. For each cluster, matrix factorization (MF) is used to predict a user’s preference transitions.

Once the top- K_1 most likely categories of a user’s next check-in are predicted, we apply another matrix factorization to predict a user’s preference on locations in the K_1 categories. By studying check-in data from a real LBSN, we found that the displacement from a user’s current location to the next one follows power-law distribution. We thus combine such geographical influence and users’ preference on locations to make the top- K_2 POI recommendations in each category, producing top- $(K_1 * K_2)$ POIs for a target user given her current check-in information.

The contributions of this work are as follows: (1) To have a better understanding of a user’s preference transitions, we propose the concept of *story* to structure her check-in activities. In a story, two consecutive check-in activities happened within a predefined time interval, thus are better correlated with each other. (2) Due to data sparsity of individual users, we propose to cluster similar users for accurate POI recommendation, taking into account both category information and temporal effects. (3) For each user cluster, we apply MF to predict the possible categories of a user’s next check-in location, given the category of her current location. (4) Once the possible categories of next check-in locations are inferred, we again apply MF to predict a user’s preference on the locations in the corresponding categories, considering geographical influence. (5) We evaluate the performance of the proposed POI recommendation method on real data collected from Gowalla. Experimental results show that our approach significantly outperforms the state-of-the-art solutions by at least 39.75% in terms of recall.

2. RELATED WORK

Motivated by its success in traditional recommendation scenarios (e.g., movie recommendation), collaborative filtering has been widely applied to POI recommendation. In [10], the geographical influence was captured by assuming a power-law distribution of a user’s check-ins. A fused recommendation framework was proposed to incorporate the user preference, geographical and social influence into one recommendation process. Moreover, the work [1] studied users’ multi-center check-in behavior and proposed a Multi-center Gaussian Model (MGM) to capture the geographical influence. Then, a matrix factorization model was fused with social and geographical influence for POI recommendation. Additionally, POI recommendation has also been collectively studied with other recommendation tasks (e.g., activity recommendation). For instance, Zheng et al. [11] proposed three collective tensor and matrix factorization models to provide personalized POI and activity recommendation.

Modeling users’ mobility using Markov model is another popular method for POI recommendation. In [7], an hybrid method for predicting users’ mobility using Hidden Markov Models (HMMs) was proposed. Users’ location histories were clustered according to their characteristics, and for

each cluster, a HMM was trained, treating location characteristics as unobservable parameters. In [2], a mobility Markov chain was built for individual users to keep track of a set of previous locations visited. Each state corresponds to a frequent POI (ranked by decreasing order of importance). The next location was predicted based on observations of a user’s mobility behavior over certain time periods and the recent locations that she has visited.

Most existing Markov model based approaches are evaluated using GPS datasets which are typically of small sizes, but are very dense. On the other hand, LBSNs have millions of users but most users’ location data is sparse. So Markov model based approaches may not produce accurate POI recommendation in LBSNs (as demonstrated in Section 4). Although our approach also assumes Markov property, we apply matrix factorization to handle large-scale POI recommendation with sparse input data.

Besides collaborative filtering and markov model, other machine learning techniques are also applied to recommend POI. In [8], the authors identified a set of relevant features from three aspects: (1) user mobility features, including historical visits, social influence, etc.; (2) global mobility features, including total number of check-ins at a location, geographical distance, etc.; (3) temporal features, including the number of check-ins at a location of certain category in a given hour, etc. Linear ridge regression and M5 decision trees, were applied to generate POI recommendation.

3. POI RECOMMENDATION APPROACH

We denote a set of LBSN users by $\mathcal{U} = \{u_1, u_2, \dots\}$, a set of locations by $\mathcal{L} = \{l_1, l_2, \dots\}$. C denotes the category tree, C^l denotes the vector of categories at level l , and $c_i^l \in C^l$ denotes the i^{th} category at level l (see Fig. 1 as an example).




Figure 1: Category tree.

In order to better understand users’ preference transition patterns, different from previous work (e.g., [4]) that studied users’ check-in behavior in a whole day, we split a user’s check-in history into several non-overlapping sequences according to the time interval between two consecutive check-ins. Specifically, if two consecutive check-in activities happened within a time interval Δt (e.g., 6 hours¹), they are assigned to the same sequence, which is defined as a *story*. Note that Δt is used to determine if two *consecutive* check-ins are clustered, but is not used to constrain the length of a story, i.e., the time difference between the first and last check-in may be longer than Δt . Fig. 2 shows an example of a user’s stories. It is believable that the consecutive check-in activities in a story are better correlated with each other (semantically meaningful) than those in different stories. For instance, a check-in at pub often follows a check-in at restaurant in the evening. On the other hand, a check-in at bus stop in the early morning and a check-in

¹The parameter Δt is application-dependent and can be chosen by analyzing user’s behavior. In our experiments, 6-hour is a suitable value to achieve accurate recommendation.



Figure 2: An example of a user’s stories.

at gym in the late afternoon, although are consecutive, are less correlated. In our approach, we focus on studying preference transitions happened within stories, by which users’ preference transitions over location categories can be more meaningfully predicted, thus increasing the accuracy of POI recommendation.

3.1 User clustering

By leveraging other similar users’ preference transition patterns, collaborative filtering is a promising method to improve the accuracy of preference transition prediction. Therefore, it is important to cluster users based on the similarity of their check-in behavior (preference), so as to make collaborative filtering based predictions more accurate [6].

Firstly, we identify a set of features to describe a user’s check-in behavior from the temporal distribution of her check-in frequency in each location category. The temporal information is considered from two aspects: the hours of the day and the days of the week. Specifically, we partition a day (24 hours) into working hours (from 08:00am to 17:59pm) and leisure time (from 18:00pm to 7:59am of the next day), and divide one week (7 days) into weekdays (from Mondays to Fridays) and weekends (Saturdays and Sundays). For each temporal period, we derive a user’s check-in frequency in each location category. If there are $|C^l|$ (l indicates the level of category) categories, we finally have $4 \times |C^l|$ check-in frequency, which can be used as features to characterize her check-in behavior.

Secondly, we apply k -means clustering algorithm to cluster users into G groups $\{\mathcal{U}_g\}_{g=1}^G$. The k -means clustering algorithm is applied due to its simplicity (i.e., suitable for large-scale LBSN with millions of users). More sophisticated clustering algorithms (e.g., expectation-maximization algorithm) can also be applied, but the discussion on the tradeoff between the improved clustering quality and the increased computational complexity is beyond the scope of this work.

3.2 Preference transition prediction

For each user cluster $\mathcal{U}_g \subset \mathcal{U}$, we apply matrix factorization, the state-of-the-art collaborative filtering technique for recommendation [3, 5], to predict the probability that a user would follow a given preference transition, based on the category of current POI. The input data to matrix factorization is a transition probability matrix $P \in \mathcal{R}^{|\mathcal{U}_g| \times |C^l|^2}$, which represents the probabilities of users’ preference transitions, where $|C^l|$ denotes the number of categories at category level l . These probabilities are derived from users’ past check-in activities in corresponding location categories. By applying matrix factorization, the transition probability matrix P is factorized into one user specific matrix $U \in \mathcal{R}^{k \times |\mathcal{U}_g|}$ and one preference transition specific matrix $V \in \mathcal{R}^{k \times |C^l|^2}$, where k is the dimensionality of the latent factor vector (each column of U or V) that characterizes the corresponding user or preference transition:

$$P \approx U^T V. \quad (1)$$

For a user u , her latent factor vector (i.e., U_u) measures the extent of u ’s preference on transitions which have high values on the corresponding latent factors; for a preference transition v (e.g., $c_a^l \rightarrow c_b^l$), its latent factor vector (i.e., V_v) measures the correlation between v and the corresponding latent factors. Therefore, the resulting $U_u^T V_v$ captures the correlation between user u and preference transition v (i.e., the predicted probability of u ’s preference on transition v). In order to accurately approximate the probabilities that users would follow certain preference transitions, the objective function of matrix factorization is defined as follows:

$$\min_{U, V} \sum_{u=1}^{|\mathcal{U}_g|} \sum_{v=1}^{|C^l|^2} I_{uv} (P_{uv} - U_u^T V_v)^2 + \lambda (\|U\|_F^2 + \|V\|_F^2), \quad (2)$$

where I_{uv} is 1 if user u has preference transition v , and 0 otherwise. $\lambda (\|U\|_F^2 + \|V\|_F^2)$ is a regularization term to avoid overfitting. $\|\cdot\|_F$ denotes the Frobenius norm.

Once the MF model is built, given the target user u and her current POI’s category c_a^l , we predict the probability that she will check-in at a POI with category c_b^l as:

$$p_u(c_b^l | c_a^l) \propto U_u^T V_v, \quad (3)$$

where v denotes the preference transition $c_a^l \rightarrow c_b^l$. $|C^l|$ predictions are obtained and we rank these location categories in descending order based on the predicted probabilities. We choose top- K_1 most likely categories for personalized POI recommendation that will be elaborated in the next section.

3.3 Category-aware POI recommendation

In order to recommend a list of POIs that are likely to interest users, we consider two factors that influence users’ check-in behavior. The first one is users’ preference on POIs. By inferring possible categories of the next check-in POI, the prediction of a user’s preference becomes more focused (i.e., only POIs in the corresponding categories are considered), which means the preference can be more reliably predicted. The second one is geographical influence. For instance, although a user often visits a Japanese restaurant, if her current location is far away from that restaurant, she may choose a nearby restaurant instead.

Previous study about users’ check-in behavior in LBSNs has demonstrated that a considerable fraction of users’ check-ins are generated at POIs which were never visited before [9]. Similar to the preference transition prediction, we again apply matrix factorization to predict a user’s preference on a given POI, taking into account other *similar* users’ preference. The input matrix $P' \in \mathcal{R}^{|\mathcal{U}_g| \times |\mathcal{L}|}$ represents users’ preference on POIs where the preference is derived as the corresponding user’s normalized check-in frequency ($\in [0, 1]$) at a certain POI. Eq. 2 can be reused to factorize P' into one user specific matrix $\tilde{U} \in \mathcal{R}^{k \times |\mathcal{U}_g|}$ and one POI specific matrix $O \in \mathcal{R}^{k \times |\mathcal{L}|}$, such that $P' \approx \tilde{U}^T O$, where k is the dimensionality of the latent factor vectors that characterize users and POIs. The probability that the target user u likes a given POI o is determined by:

$$p_{uo}^p \propto \tilde{U}_u^T O_o, \quad (4)$$

where $\tilde{U}_u \in \mathcal{R}^{k \times 1}$ and $O_o \in \mathcal{R}^{k \times 1}$ are the latent factor vectors of the target user u and the target POI o respectively.

By studying the real check-in data, we found users’ displacements (distance between two consecutive check-in loca-

tions) demonstrate power-law property²: users prefer visiting nearby venues but traveling to farther POIs is also likely. Based on this observation, we model a user’s check-in probabilities to displacements using a power-law distribution:

$$p^d = a \times D^{-b}, \quad (5)$$

where p^d is the probability that the user will check-in at a POI, and D is the displacement between the user’s current POI and that POI. a (a normalizing constant) and b (a decay exponent) are power-law distribution parameters. In practice, a lower bound D_{min} is needed such that values larger than D_{min} follow power-law property. After calculating the normalizing constant, Eq. 5 can be rewritten as:

$$p^d = \frac{b-1}{D_{min}} \left(\frac{D}{D_{min}} \right)^{-b}. \quad (6)$$

By applying maximum likelihood method, the decay exponent b can be estimated as follows:

$$\hat{b} = 1 + n \left[\sum_{i=1}^n \ln \frac{D_i}{D_{min}} \right]^{-1}, \quad (7)$$

where D_i ($\geq D_{min}$) are observed displacements of the user.

By inferring users’ preference on POIs and their personalized mobility patterns (with power-law property), we fuse these two factors to derive the probability that the target user u will check-in at a POI o :

$$p_{uo} = p_{uo}^p p_{uo}^d, \quad (8)$$

where p_{uo}^p is calculated by Eq. 4 and p_{uo}^d is calculated by Eq. 6. Note that only POIs in the corresponding categories (predicted in Section 3.2) are considered.

In preference transition prediction stage, we chose top- K_1 possible categories of the next check-in POI. For each category, we sort relevant POIs according to the probabilities that they will be visited (by Eq. 8). We then choose top- K_2 POIs with the highest probabilities for each category to construct a recommendation list of totally $K_1 * K_2$ POIs for the target user.

4. EVALUATION

4.1 Methodology

The experimental data prepared for this study was collected from Gowalla, a popular LBSN, which has more than 600,000 users since November 2010 and was acquired by Facebook in December 2011. We collected users’ check-in data (longitude, latitude, timestamp, categories, etc.) generated before June 1, 2011 in the following cities: Austin, Chicago, Houston, Los Angeles and San Francisco. The statistics of the data are summarized in Table 1, in which N_u denotes the number of users, N_l denotes the number of locations, N_c denotes the number of check-ins, \tilde{N}_u denotes the average number of visited users for a location, and \tilde{N}_l is used to denote the average number of visited locations for a user. Moreover, we have also collected the category of each observed location. Locations in Gowalla are grouped into 7 main categories, i.e., Community, Entertainment, Food, Nightlife, Outdoors, Shopping and Travel. In each main category, the locations are classified into different subcategories (there are totally 134 categories at the second level).

Table 1: The statistics of experimental data

	N_u	N_l	N_c	\tilde{N}_u	\tilde{N}_l
Austin	24,070	51,118	1,935,677	21.95	46.62
Chicago	13,845	37,050	486,558	7.53	20.16
Houston	11,138	29,383	512,977	9.89	26.08
L.A.	21,633	75,301	1,296,953	10.06	35.02
S.F.	21,585	64,758	1,542,133	13.76	41.29

For each user, we sort her check-ins in chronological order and select the first 70% of check-ins to construct the training data to predict the remaining 30% check-ins. For each algorithm in comparison, we run ten times and report the average results.

We compare our approach with the following state-of-the-art recommendation models: (1) **BasicMF**. To start with, we use a basic MF model as the baseline (prediction using Eq. 1). This model only considers the user preference for POI recommendation, without taking into account other auxiliary information (e.g., geographical, social, and temporal information). (2) **GeoCF**. In this approach [10], the geographical influence is considered (by assuming a power-law distribution) and integrated with a user-based collaborative filtering algorithm. A unified POI recommendation framework is proposed to linearly combine the user preference and geographical influence. (3) **MGM**. This approach [1] applies Multi-center Gaussian Model (MGM) to study users’ multi-center check-in behavior. This is based on the observations that users tend to check-in around several centers, where the check-in probability follows a Gaussian distribution at each center. A fusion framework is proposed to fuse the user preference and MGM based check-in probability to recommend POIs. (4) **Markov**. In [2], by applying mobility Markov chain (each state corresponds to a frequent POI), a user’s next location is predicted based on the observations of her mobility behavior over certain temporal period and the recent locations that she has visited. (5) **ML**. This approach [8] identifies a set of features (e.g., user mobility features, global mobility features, temporal features) to describe users’ check-in behavior. M5 decision tree is applied to predict the next POIs of a user.

To measure the accuracy of POI recommendation, we use two standard metrics: Precision@ K and Recall@ K [10, 1]. Specifically, we say a POI is successfully predicted if it is in the recommendation list. Precision@ K defines the ratio of successfully predicted POIs to the K recommendations, and Recall@ K defines the ratio of successfully predicted POIs to the number of POIs to be predicted (i.e., the size of test set). We will demonstrate how the performance of recommendation models varies with different values of K .

4.2 Results

4.2.1 Different components of our approach

We first study how the precision and recall of our preference transition prediction varies with different sizes of the recommendation list (i.e., top- K_1). From Fig. 3 we observe the larger the recommendation list size, the higher the recall is. This is because larger recommendation list size increases the probability that a user’s next check-in is hit by one of the recommended POIs. For example, the recall reaches 0.9 (even higher) at level 1 categories, when K_1 is set to 5. At level 2 categories, the lowest recall is still higher than 0.4

²This observation is consistent with previous work [1, 10].




Figure 3: Impacts of recommendation list size (top- K_1 recommendation).




Figure 4: Impacts of recommendation list size (top- $K_1 * K_2$ recommendation).

in San Francisco data, and in Houston data, it is as high as 0.6. This demonstrates that users’ preference transition can be accurately predicted, which prepares a good basis for the personalized POI recommendation at the next stage.

Table 2: Performance of different components (top-2 categories, top-3 POIs)

	Category level 1		Category level 2	
	Recall	Precision	Recall	Precision
U Preference	0.1907	0.0318	0.2320	0.0387
Geo Influence	0.0719	0.0120	0.0892	0.0149
Both factors	0.2354	0.0392	0.2955	0.0492

Once possible categories of the next POI are predicted, a list of POIs (in the corresponding categories) are recommended by considering two factors: user preference and geographical influence. In Tab. 2, we report the accuracy of our approach with respect to (1) considering only the user preference, (2) considering only the geographical influence, and (3) considering both the user preference and geographical influence. Note that we select top-2 most likely categories (preference transition prediction) and for each category we recommend top-3 most likely POIs. Obviously, combining the two factors achieves the best performance. When only one factor is considered, we observe that the user preference plays a more important role: with respect to the scenario where only the geographical influence is used, the precision and recall is improved by around 165% at category level 1 and around 160% at category level 2.

Finally, we study the impacts of recommendation list size on both preference transition prediction (top- K_1 categories) and category-aware POI recommendation (top- K_2 POIs). We choose the values of K_1 and K_2 from 1 to 5, and form 5*5 combinations in total. Fig. 4 shows the recall and precision with respect to different combinations of K_1 and K_2 at both category level 1 and level 2. As expected, the general trend is that the larger the value of $K_1 * K_2$, the higher the recall and precision are. We observe that the advantage of considering a certain category level for POI recommendation varies with different $K_1 * K_2$ combinations. For instance, when $K_1 = 3$ and $K_2 = 3$ are applied, POI recommendation at category level 2 is more accurate (higher recall) than that at category level 1; while when $K_1 = 3$ and $K_2 = 5$, POI

recommendation at category level 1 is more accurate. We also observe when the value of $K_1 * K_2$ is fixed, different values of K_1 and K_2 (e.g., 3*5 versus 5*3) generates slightly different performance. So in the following experiments, we choose the optimal settings of the category level and $K_1 * K_2$ combination to maximize the recommendation accuracy.

4.2.2 Comparison study

Tab. 3 summarizes the performance of all recommendation models in individual cities. Top-6 POIs are recommended by each model. In particular, for our approach, category level 2 is considered; K_1 and K_2 are set to 2 and 3 respectively. Surprisingly, although it is one of the most accurate algorithms for recommendation tasks, BasicMF incurs the lowest recall and precision. This is because MF is effective when users provide explicit ratings (e.g., on 5-point likert scale), but for POI recommendation, only check-in frequency is provided, which cannot accurately reflect a user’s real interests in locations. The Markov model achieves better performance than BasicMF, but it is still much less accurate than other methods. The reason is that Markov models are typically designed for the mobile scenario where individual’s mobility data is dense (e.g., GPS location data), such that a reliable Markov predictive model can be built. However, in LBSNs, users’ location data is very sparse, thus greatly degrading the accuracy of Markov model. In contrast, GeoCF, MGMMF and ML are particularly designed for LBSNs. By integrating the geographical influence into collaborative filtering, GeoCF and MGMMF outperform ML, demonstrating the benefits of combining the geographical information and user preference for POI recommendation. Although ML also considers the geographical influence, it is used as a single feature, and is combined with other features using supervised learning algorithms, without deeply exploiting the correlation with other factors. In all cities, our approach consistently and significantly outperforms all other models, demonstrating the advantage of our two-stage approach. We also observe that the accuracy of recommendation in different cities varies a lot. For example, the recall@6 in Chicago data is 85.55% higher than that in Austin data. This is probably due to the different user populations and check-in behavior patterns in different cities.

Table 3: Performance comparison

	Houston		Chicago		Los Angeles		San Francisco		Austin	
	Recall	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall	Precision
Our approach	0.2785	0.0464	0.2915	0.0485	0.2027	0.0338	0.1590	0.0265	0.1571	0.0262
BasicMF	0.0823	0.0137	0.0855	0.0143	0.0724	0.0121	0.0655	0.0109	0.0639	0.0107
GeoCF	0.1645	0.0274	0.2270	0.0378	0.1642	0.0274	0.1260	0.0210	0.1217	0.0203
MGMMF	0.1587	0.0265	0.1855	0.0309	0.1524	0.0254	0.1122	0.0187	0.1174	0.0196
Markov	0.1024	0.0171	0.1157	0.0193	0.0957	0.0160	0.0877	0.0146	0.0855	0.0143
ML	0.1524	0.0254	0.1698	0.0283	0.1323	0.0221	0.1106	0.0184	0.1157	0.0193




Figure 5: Performance comparison with varying recommendation list size.

Finally, we show the performance of all recommendation models with varying recommendation list size in Fig. 5. Note that the results are averaged across all cities. Obviously, for all models, larger recommendation list size increases the recall but decreases the precision, which is consistent with previous experimental results (see Fig. 3 and 4). The maximum recommendation list size is set to 20, as larger size may affect the effective interactions between a user and the LBSN via mobile devices. Again our approach outperforms all other recommendation models when different list sizes are applied. In summary, our approach outperforms GeoCF, MGMMF, Markov and ML by 39.75%, 54.40%, 119.16%, and 78.72% respectively, in terms of recall.

5. CONCLUSION

In this work, we decompose the process of POI recommendation into two stages: (1) we predict a user's preference transition (over POI categories) based on her check-in behavior. (2) Once possible categories of the next POI are predicted, a set of POIs in the predicted categories are recommended, taking into account both the user preference and geographical influence. Our two-stage POI recommendation approach improves the accuracy of POI recommendation via deeply understanding users' preference transition patterns and shrinking the POI candidate list size. Extensive experiments on real data demonstrate that our approach significantly outperforms other state-of-the-art POI recommendation models in different settings. As for the future work, we plan to incorporate users' textual comments (e.g., tips in Foursquare) to further improve the prediction accuracy of users' preference on POIs.

6. ACKNOWLEDGEMENT

This work was partially supported by (1) the grant *Reconcile: Robust Online Credibility Evaluation of Web Content* from Switzerland through the Swiss Contribution to the enlarged European Union, and (2) the Interactive and Digital Media Programme Office (IDMPO), National Research Foundation (NRF) hosted at Media Development Authority (MDA) under Grant No.: MDA/IDM/2012/8/8-2 VOL 01.

7. REFERENCES

- [1] C. Cheng, H. Yang, I. King, and M. R. Lyu. Fused matrix factorization with geographical and social influence in location-based social networks. In *Proceedings of the 26th AAAI*, 2012.
- [2] S. Gamba, M.-O. Killijian, and M. N. del Prado Cortez. Next place prediction using mobility markov chains. In *Proceedings of the First Workshop on Measurement, Privacy, and Mobility*, 2012.
- [3] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [4] W. Li, C. Eickhoff, and A. P. de Vries. Want a coffee?: predicting users' trails. In *Proceedings of the 35th ACM SIGIR*, 2012.
- [5] X. Liu and K. Aberer. Soco: A social network aided context-aware recommender system. In *Proceedings of the 22nd WWW*, 2013.
- [6] H. Ma, H. Cao, Q. Yang, E. Chen, and J. Tian. A habit mining approach for discovering similar mobile users. In *Proceedings of the 21st WWW*, 2012.
- [7] W. Mathew, R. Raposo, and B. Martins. Predicting future locations with hidden markov models. In *Proceedings of ACM UbiComp*, 2012.
- [8] A. Noulas, S. Scellato, N. Lathia, and C. Mascolo. Mining user mobility features for next place prediction in location-based services. In *Proceedings of the 12th IEEE ICDM*, 2012.
- [9] A. Noulas, S. Scellato, N. Lathia, and C. Mascolo. A random walk around the city: New venue recommendation in location-based social networks. In *Proceedings of IEEE SocialCom*, 2012.
- [10] M. Ye, P. Yin, W.-C. Lee, and D.-L. Lee. Exploiting geographical influence for collaborative point-of-interest recommendation. In *Proceedings of the 34th ACM SIGIR*, 2011.
- [11] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang. Towards mobile intelligence: Learning from gps history data for collaborative recommendation. *Artif. Intell.*, 184-185:17–37, 2012.