Towards Personalized Context-Aware Recommendation by Mining Context Logs through Topic Models

Kuifei Yu[1,2], Baoxian Zhang[1], Hengshu Zhu[2,3], Huanhuan Cao[2], and Jilei Tian[2]

[1] Graduate University of Chinese Academy of Sciences
[2] Nokia Research Center
[3] University of Science and Technology of China
{kuifei.yu,happia.cao,jilei.tian}@nokia.com, bxzhang@gucas.ac.cn, zhs@ustc.edu.cn

Abstract. The increasing popularity of smart mobile devices and their more and more powerful sensing ability make it possible to capture rich contextual information and personal context-aware preferences of mobile users by user context logs in devices. By leveraging such information, many context-aware services can be provided for mobile users such as personalized context-aware recommendation. However, to the best knowledge of ours, how to mine user context logs for personalized context-aware recommendation is still under-explored. A critical challenge of this problem is that individual user's historical context logs may be too few to mine their context-aware preferences. To this end, in this paper we propose to mine common context-aware preferences from many users' context logs through topic models and represent each user's personal context-aware preferences as a distribution of the mined common context-aware preferences. The experiments on a real-world data set contains 443 mobile users' historical context data and activity records clearly show the approach is effective and outperform baselines in terms of personalized context-aware recommendation.

Keywords: Personalization, Recommender System, Context-Aware, Mobile Users, Latent Dirichlet Allocation (LDA).

1 Introduction

Recent years have witnessed the increasing popularity of smart mobile devices, such as smart phones and pads. These devices are usually equipped with multiple context sensors, such as GPS sensors, 3D accelerometers and optical sensors, which enables them to capture rich contextual information of mobile users and thus support a wide range of context-aware services, including context-aware tour guide [15], location based reminder [13] and context-aware recommendation [2,9,16,10], etc. Moreover, these contextual information and users' corresponding activity (e.g., browsing web sites, playing games and chatting by Social Network Services) can be recorded into *context logs* to be used for mining

P.-N. Tan et al. (Eds.): PAKDD 2012, Part I, LNAI 7301, pp. 431–443, 2012.

users' personal context-aware preferences. By considering both context-aware preferences and the current contexts of users, we may be able to make more personalized context-aware recommendations for mobile users.

Indeed, the personalized context-aware recommendations can provide better user experiences than general context-aware recommendations which only take into account contexts but not users' different personal preferences under same contexts. In recent years, many researchers studied the problem of personalized context-aware recommendation [17,12,9]. However, most of this work is based on item ratings generated by users, which are difficult to obtain in practise. In contrast, user activity records in context logs are much easier to get for mobile users.

To the best knowledge of ours, how to mine personal context-aware preferences from context logs and then make personalized context-aware recommendations is still under-explored. To this end, in this paper we attempt to leverage mining user context logs for personalized context-aware recommendation. However, a critical challenge of the problem is that individual user's context logs usually have no sufficient training data for mining personal context-aware preferences. To be specific, as showed in Table 1, it can be observed that many context records have no corresponding activity record. As a result, if we only leverage individual user's context logs for context-aware preference mining, it will be very difficult to learn personal context preferences for recommendation, which is also reflected by our experiments on a real world data set. To address this problem, in this paper, we propose to mine **C**ommon **C**ontext-aware **P**references (CCPs) from many users' context logs through topic models and represent each user's personal context-aware preferences as a distribution of the mined common context-aware preferences. To be specific, first we extract bags of *Atomic Context-Aware Preference (ACP) Features* for each user from their historical context logs. Then, we propose to mine CCPs from users' ACP-feature bags through topic models. Finally, we make recommendations according to the given contexts and CCP distributions of users. Figure 1 illustrates our procedure for generating personalized context-aware recommendation. In addition, we evaluate our proposed approach in a real-world data set of context logs collected from 443 mobile phone users spanning for several months, which contains more than 8.8 million context records, 665 different interactions in 12 content categories.

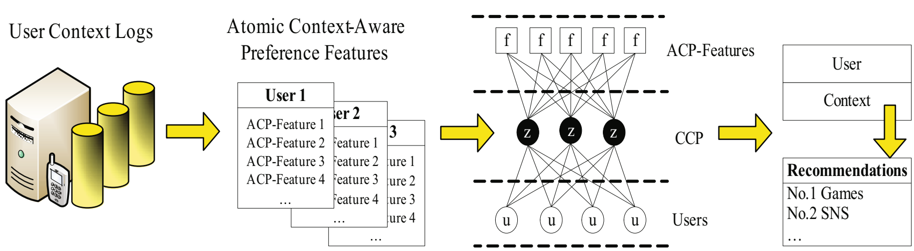

Fig. 1. The procedure of personalized context-aware recommendation for mobile users

Table 1. A toy context log from real-world data set

Timestamp	Context	Activity record
t_1	{(Day name: Monday),(Time range: AM8:00-9:00)), (Profile: General),(Location: Home)}	Null
t_2	{(Day name: Monday),(Time range: AM8:00-9:00)), (Profile: General),(Location: On the way)}	Play action games
t_3	{(Day name: Monday),(Time range: AM8:00-9:00)), (Profile: General),(Location: On the way)}	Null
	......	
t_{359}	{(Day name: Monday),(Time range: AM10:00-11:00), (Profile: Meeting),(Location: Work place)}	Null
t_{360}	{(Day name: Monday),(Time range: AM10:00-11:00), (Profile: Meeting),(Location: Work place)}	Browsing sports web sites
	......	
t_{448}	{(Day name: Monday),(Time range: AM11:00-12:00), (Profile: General),(Location: Work place)}	Play with Social Network Serivce
t_{449}	{(Day name: Monday),(Time range: AM11:00-12:00), (Profile: General),(Location: Work place)}	Null

The results clearly demonstrate the effectiveness of the proposed approach and indicate some inspiring conclusions.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of related works. Then, Section 3 presents the idea of making personalized context-aware recommendation by mining context logs for mining users' context-aware preferences, and Section 4 presents how to mine common context-aware preferences through topic models. Section 5 reports our experimental results on a real world data set. Finally, in Section 6, we conclude this paper.

2 Related Work

Today, the powerful sensing abilities of smart mobile devices enable to capture the rich contextual information of mobile users, such as location, user activity, audio level, and so on. Consequently, how to leverage such rich contextual information for personalized context-aware recommendation has become a hot problem which dramatically attracts many researchers' attention.

Many previous works about personalized context-ware recommendation for mobile users have been reported. For example, Tung *et al.* [14] have proposed a prototype design for building a personalized recommender system to recommend travel related information according to users' contextual information. Park *et al.* [12] proposed a location-based personalized recommender system, which can reflect users' personal preferences by modeling user contextual information through Bayesian Networks. Bader *et al.* [2] have proposed a novel context-aware approach to recommending points-of-interest (POI) for users in an automotive scenario. Specifically, they studied the scenario of recommending gas stations for car drivers by leveraging a Multi-Criteria Decision Making (MCDM) based methods to modeling context and different routes. However, most of these works only leverage individual user's historical context data for modeling personal context-aware preferences, and do not take into account the problem of insufficient personal training data.

Actually, the problem of insufficient personal training data is common in practice and many researchers have studied how to address this problem. For example, Woerndl *et al.* [16] proposed a hybrid framework named "play.tools" for recommending mobile applications by leveraging users' context information. This recommendation framework are based on what other users have installed in similar context will be liked by a given user. Kim *et al.* [9] investigated several Collaborative Filtering (CF) based approaches for recommendation and developed a memory based CF approach to providing context-aware advertisement recommendation. Specially, the proposed approach can leverage a classification rule of decision tree to understand users' personal preference. Zheng *et al.* [17] have studied a model based CF approach to recommending user locations and activities according to users' GPS trajectories. The approach can model user, location and activity as a 3-dimensional matrix, namely tensor, and perform tensor factorization with several constraints to capture users' preferences. Alexandros *et al* [10] proposed a model based CF approach for making recommendation with respect to rich contextual information, namely multiverse recommendation. Specifically, they modeled the rich contextual information with item by N-dimensional tensor, and proposed a novel algorithm to make tensor factorization. In a word, most of these approaches are based on rating logs of mobile users and the objective is to predict accurate ratings for the unobserved items under different contexts. However, usually we cannot obtain such rating data in user mobile devices. In contrast, it is easier to collect context logs which contain users' historical context data and activity records, which motivates our work for exploring how to leverage context logs for personalized context-aware recommendation.

The proposed approach in this paper exploits topic models for learning users' CCPs. Indeed, topic models are widely used in text retrieval and information extraction. Typical topic models include the Mixture Unigram (MU) [11], the Probabilistic Latent Semantic Indexing (PLSI) [8], and the Latent Dirichlet Allocation (LDA) [4]. Most of other topic models are extended from the above ones for satisfying some specific requirements. In our approach, we exploit the widely used LDA model.

3 Preliminary

As mentioned in Section 1, smart devices can capture the historical context data and corresponding activity records of users through multiple sensors and record them in context logs. For example, Table 1 shows a toy context log of a mobile user, which contains several *context records*, and each context record consists of a timestamp, the most detailed available context at that time, and the corresponding user activity record captured by devices. A context consists of several *contextual features* (e.g., Day name, Time range, and Location) and their corresponding values (e.g., Saturday, AM8:00-9:00, and Home), which can be annotated as *contextual feature-value pairs*. And we mention "available" because a context record may miss some context data though which context data which should be collected is usually predefined. For example, the GPS coordinate is not available when the user is indoor. Moreover, interaction records can be empty (denoted as "Null") because the user activities which can be captured by devices do not always happen.

It is worth noting that we transform raw location based context data such as GPS coordinates or cell Ids into social locations which have explicit meanings such as "Home" and "Work place" by some existing location mining approaches (e.g., [5]). The basic idea of these approaches is to find clusters of user location data and recognize their social meaning by time pattern analysis. Moreover, we also manually transform the raw activity records to more general ones by mapping the activity of using a particular application or playing a particular game to an activity category. For example, we can transform two raw activity records "*Play Angry Birds*" and "*Play Fruit Ninja*" to same activity records "*Play action games*". In this way, the context data and activity records in context logs are normalized and the data sparseness is some how alleviated for easing context-aware preference mining.

Given a context $C = \{p\}$ where p denotes an atomic context, i.e., a contextual feature-value pair, the probability that a user u prefers activity a can be represented as

$$P(a|C,u) = \frac{P(a,C|u)P(u)}{P(C,u)} \propto P(a,C|u) \propto \prod_p P(a,p|u),$$

where we assume that the atomic contexts are mutually conditionally independent given u.

Then the problem becomes how to calculate $P(a,p|u)$. According to our procedure, we introduce a variable of CCP denoted as z, and thus we have

$$P(a,p|u) = \sum_z P(a,p,z|u) \propto \sum_z P(a,p|z,u)P(z,u) \propto \sum_z P(a,p|z)P(z|u),$$

where we assume that a user's preference under a context only relies on the CCPs and his (her) context-aware preferences in the form of their distribution on the CCPs, rather than other information of the user. Therefore, the problem is further converted into learning $P(a,p|z)$ and $P(z|u)$ from many users' context logs, which can be solved by widely used topic models. In the next section, we present how to utilize topic models for mining CCPs, i.e., $P(a,p|z)$, and accordingly make personalized context-aware recommendation.

4 Mining Common Context-Aware Preferences through Topic Models

Topic models are generative models that are successfully used for document modeling. They assume that there exist several topics for a corpus D and a document d_i in D can be taken as a bag of words $\{w_{i,j}\}$ which are generated by these topics. For simplicity, we refer the co-occurrence of a user activity a

and the corresponding contextual feature-value pair p, i.e., (a, p), as *Atomic Context-aware Preference feature*, and *ACP-feature* for short. Intuitively, if we take ACP-features as words, take context logs as bags of ACP-features to correspond documents, and take CCPs as topics, we can take advantage of topic models to learn CCPs from many users' context logs.

However, raw context logs are not naturally in the form of bag of ACP-features so we need some preprocessing for extracting training data. Specially, we first remove all context records without any activity record and then extract ACP-feature from the remaining ones. Given a context record $< Tid, C, a >$ where Tid denotes the timestamp, $C = \{p_1, p_2, ..., p_l\}$ denotes the context and a denotes the activity, we can extract l ACP-features, namely, (a, p_1), (a, p_2), ..., (a, p_l). For simplicity, we refer the bag of ACP-features extracted from user u's context log as the *ACP-feature bag* of u.

Among several existing topic models, in this paper, we leverage the widely used Latent Dirichlet Allocation model (LDA) [4]. According to LDA model, the ACP-feature bag of user u_i denoted as d_i is generated as follows. First, before generating any ACP-feature bag, K prior ACP-feature conditional distributions given context-aware preferences $\{\phi_z\}$ are generated from a prior Dirichlet distribution β. Secondly, a prior context-aware preference distribution θ_i is generated from a prior Dirichlet distribution α for each user u_i. Then, for generating the j-th ACP-feature in d_i denoted as $w_{i,j}$, the model firstly generates a CCP z from θ_i and then generates $w_{i,j}$ from ϕ_z. Figure 2 shows the graphic representation of modeling ACP-feature bags by LDA.

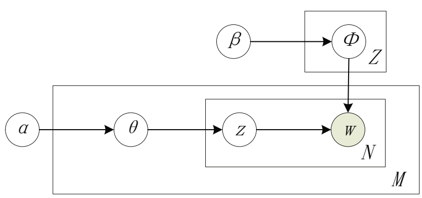

Fig. 2. The graphical model of LDA

In our approach, the objective of LDA model training is to learn proper estimations for latent variables θ and ϕ to maximize the posterior distribution of the observed ACP-feature bags. In this paper, we choose a Markov chain Monte Carlo method named Gibbs sampling introduced in [6] for training LDA models efficiently. This method begins with a random assignment of CCPs to ACP-features for initializing the state of Markov chain. In each of the following iterations, the method will re-estimate the conditional probability of assigning a CCP to each ACP-feature, which is conditional on the assignment of all other ACP-features. Then a new assignment of CCP to ACP-features according to those latest calculated conditional probabilities will be scored as a new state of Markov chain. Finally, after rounds of iterations, the assignment will converge, which means each ACP-feature is assigned a stable and final CCP. Eventually, we can obtain the estimated values for two distributions $\{\widetilde{p}(a,p|z)\}$ and $\{\widetilde{p}(z|u)\}$, which denote the probability that the ACP-feature (a,p) appears under the CCP z, and the probability that user u has the context-aware preference z, respectively.

$$\widetilde{p}(a,p|z) = \frac{n_{(z)}^{(a,p)} + \beta}{n_{(a,p)}^{(.)} + A\beta}, \quad \widetilde{p}(z|u) = \frac{n_{(z)}^{(u)} + \alpha}{n_{(.)}^{(u)} + Z\alpha},$$

where the $n_{(z)}^{(a,p)}$ indicates the number of times ACP-feature (a,p) has been assigned to CCP z, while $n_{(z)}^{(u)}$ indicates the number of times a ACP-feature from user u's context log that has been assigned to CCP z. The A indicates the number of ACP-features from u's context log, and Z indicates the number of CCPs.

LDA model needs a predefined parameter Z to indicate the number of CCPs. How to select an appropriate Z for LDA is an open question. In terms of guaranteeing the performance of recommendation, in this paper we utilize the method proposed by Bao et al [3] to estimate Z, and we set ζ to be 10% in our experiments accordingly. Please refer to [1] for more information.

After learning CCPs represented by distributions of ACP-features, we can predict users' preference according to their historical context-aware preferences and current contexts, i.e., $P(a,C|u)$. Then, we recommend users a rank list of different categories of contents according to the preference prediction. For example, if we predict a user u is more likely willing to play action games than listen pop music, the recommendation priority of popular action games will be higher than that of recent hot pop music.

5 Experiments

In this section, we evaluate the performance of our LDA based personalized context-aware recommendation approach, namely **P**ersonalized **C**ontext-aware **R**ecommendation with **LDA** (PCR-LDA), with several baseline methods in a real-world data set.

5.1 Data Set

The data set used in the experiments is collected from many volunteers by a major manufacturer of smart mobile devices. The data set contains context logs with rich contextual information and user activities of 443 smart phone users spanning for several months. The detailed statistics of our data set are illustrated in Table 2. From table 2 we can observe that only 12.5% context records have activities, which indicates the insufficient activity records for individual user in practice. Moreover, Table 3 shows the concrete types of context data contained in our data set. In addition, in our data set, all activities can be classified into 12

Table 2. Statistics of our data set

	Number
users	443
unique activities	665
unique context	4,391
context records	8,852,187
activity-context records[†]	1,097,189

[†] activity-context records denote the context records with non-empty user activity records.

content categories, which are *Call*, *Web*, *Multimedia*, *Management*, *Games*, *System*, *Navigation*, *Business*, *Reference*, ***S**ocial **N**etwork **S**ervice (SNS)*, *Utility* and *Others*. Specifically, in our experiments, we do not utilize the categories *Call* and *others* because their activity information is clear for making recommendations. Therefore, in our experiments we utilize 10 activity categories which contain 618 activities appear in total 408,299 activity-context records.

Table 3. The types of contextual information in our data set

Data source	Data type	Value range
Time Info	Week	{Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday}
	Is a holiday?	{Yes, No}
	Day period	{Morning(AM7:00-AM11:00), Noon(AM11:00-PM14:00), Afternoon(PM14:00-PM18:00), Evening(PM18:00-PM21:00), Night(PM21:00-Next day AM7:00)}
	Time range	{AM0:00-AM1:00, AM1:00-AM2:00, AM2:00-AM3:00, ... , PM23:00-PM24:00}
System Info	Profile type	{General, Silent, Meeting, Outdoor, Pager, Offline}
Geo Info	Location	{Home, Work Place, On the way}.

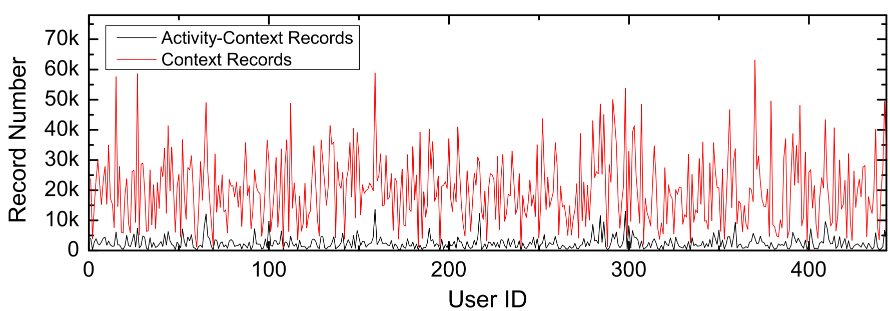

Fig. 3. The distribution of context coverage for all users

Figure 3 shows the distribution of context records and activity-context records for all users. From the figure we can see that usually though the context records of individual mobile users are sufficient, only small proportion of them have non-empty activity records and can be used as training data, which implies the limit of learning personal context-aware preferences only from individual user's context logs.

5.2 Benchmark Methods

To evaluate the recommendation performance of our approach, we chose two context-aware baseline methods as follows.

CPR stands for Context-aware Popularity based Recommendation which is a basic context-aware recommendation approach without considering personal context-aware preference. To be specific, in this approach, given a user u and a context C, we predict user preferred activities by the most frequent activities appear under C according to all users' historical context logs and recommend corresponding contents. This popularity based approach is widely used in practical recommender systems.

PCR-i stands for Personalized Context-aware Recommendation by only leveraging Individual user's context logs. To be specific, in this approach, given a user u and a context C, we rank each activity a by probability $P(a|u, C)$, which can be estimated by $P(a|u, C) \propto (\prod_{p \in C} P(a, p|u)$. The probability $P(a, p|u)$ can be calculated by $P(a, p|u) = \frac{n_{a,p}}{n_{(.)}}$, where $n_{a,p}$ and $n_{(.)}$ indicate the numbers of ACP-feature (a, p) and all ACP-features appeared in the context log of u, respectively.

5.3 Evaluation Metrics

In the experiments, we utilize 5-fold-cross validation to evaluate the performance of each recommendation approaches. To be specific, we first randomly divide each user's context log into five equal parts, then use each part as test data while use other four parts as training data for total 5-rounds of recommendation. In the test process, we only take into account the context records with non-empty activity records, and use the contexts and the content categories corresponding to the real user activity as context input and ground truth, respectively. In our experiments, each recommendation approach will return a ranked list of recommended content categories according to predicted user activities. To evaluate the performance of each approach, we leverage two different metrics as follows.

MAP@K stands for Mean Average Precision at top K recommendation results. To be specific, $MAP@K = \frac{\sum AP^{(u)}@K}{|U|}$, where $AP^{(u)}@K$ denotes the average precision at top k recommendation results on the test cases of user u, and $|U|$ indicates the number of the users. $AP^{(u)}@K$ can be computed by $\frac{1}{N_u} \sum_i \sum_{r=1}^{K} (P_i(r) \times rel_i(r))$, where N_u denotes the number of test cases for user u, r denotes a given cut-off rank, $P_i(r)$ denotes the precision on the i-th test case of u at a given cut-off rank r, and $rel_i()$ is the binary function on the relevance of a given rank.

MAR@K stands for Mean Average Recall at top K recommendation results. To be specific, $MAR@K = \frac{\sum AR^{(u)}@K}{|U|}$, where $AR^{(u)}@K$ denotes the average recall at top k recommendation results on the test cases of user u, and $|U|$ indicates the number of the users. $AR^{(u)}@K$ can be computed by $\frac{1}{N_u}\sum_i \sum_{r=1}^{K} rel_i(r)$, where N_u denotes the number of test cases for user u, r denotes a given cut-off rank, and $rel_i()$ is the binary function on the relevance of a given rank.

5.4 Overall Results of Recommendation

To evaluate our PCR-LDA recommendation approach, we compare its recommendation performance with other baselines. To be specific, according to the parameter estimation approaches introduced in Section 4, the number of CCPs for LDA training is set to be 15. In Section 5.5 we will further discuss the setting of this parameter. For the LDA training, the two parameters α and β are empirically set to be $50/Z$ and 0.2 according to discussion in [7]. Both PCR-LDA and the two baselines are implemented by C++ and the experiments are conducted on a 3GHZ×4 quad-core CPU, 3G main memory PC.

We first test the $MAP@K$ performance of each recommendation approach with respect to varying K, which are shown in Figure 4. From the results we can observe that PCR-LDA outperforms other baselines with a significant margin.

Figure 5 shows the $MAR@K$ of each recommendation approach. From the results we can observe our PCR-LDA can achieve 100% performance when $K = 10$, which means they can return recommendation list contains at least one ground

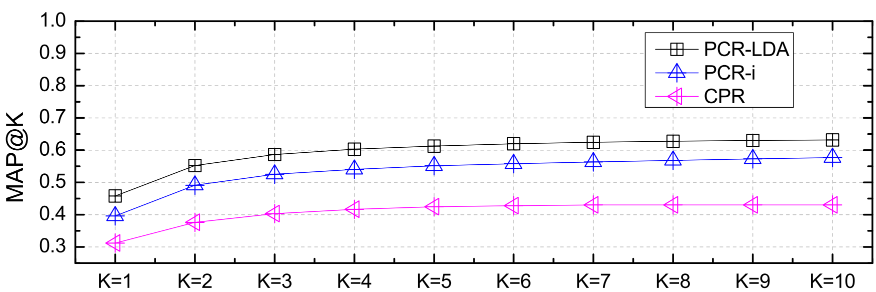

Fig. 4. The MAP@K performance of each recommendation approach

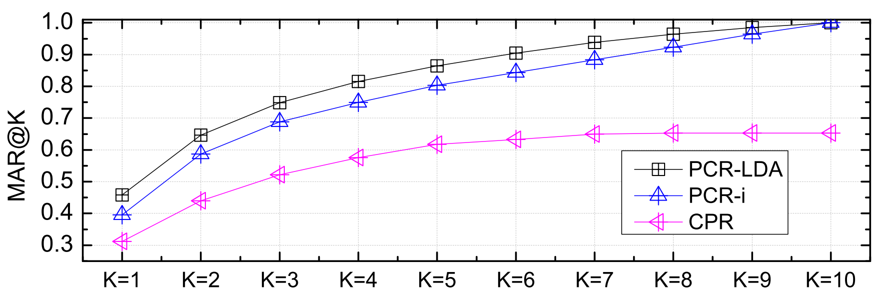

Fig. 5. The MAR@K performance of each recommendation approach

truth activities for all contexts. It is because PCR-LDA takes advantage of many users' context logs. In contrast, PCR-i has worse MAR@K due to the insufficient training data in individual user's context logs for mining context-aware preference. Moreover, due to the different context-aware preference between users, the popularity based approach CPR under-performs the other approaches.

5.5 Robustness Analysis

CPR-LDA needs a parameter Z to determine the number of CCPs. Although we can empirically select Z by estimating perplexity, we still study the impact of such parameter to our recommendation results. Figure 6 shows the MAP@10 of PCR-LDA with respect to varying Z. From the results we can observe that the MAP@10 of PCR-LDA is impacted dramatically by a relatively small Z and becomes stable with relatively big Z. It is because when a relatively small Z is selected, all ACP-features may have strong relationships with each CCP. Thus the approach is actually near to combine all users' context logs as one log for recommendation, which will introduce many noisy data. Another interesting phenomenon is that the MAP@10 peaks when Z is set to be 15, which is consistent with our experimental setting and implies the parameter selection method if effective. The experimental results of $MAP@K$ with other settings of K show the similar phenomena.

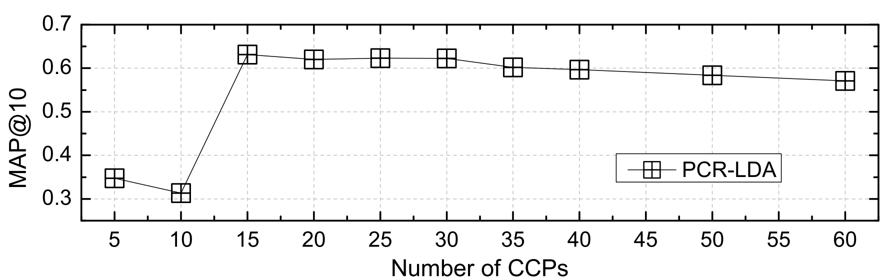

Fig. 6. The MAP@10 performance of PCR-LDA to varying number of CCPs

5.6 Case Study

In addition to the studies on the overall performance of our recommendation approach, we also study the cases in which PCR-LDA outperforms the baselines. For example, Table 4 shows the top 3 recommendation results of each approach for two test cases of two different users given the "{(Is holiday: No), (Day period: Evening), (Time range: PM22:00-23:00), (Day name: Monday), (Profile: General), (Location: Home)}", which may imply the users' leisure time at home. In this case, the activity records of user #152 and user #343's test cases are *Multimedia* and *Web*, respectively. From the results, we can observe that PCR-LDA recommend relevant content categories in the top one position. In contract, PCR-i can only recommend relevant content categories in the top one position for one test case, and the Popularity based approach CPR always recommend same content categories for all users and thus sometimes performs not well.

Table 4. An example of recommendation results for user #152 and #343

Context	{(Time range: PM22:00-23:00),(Is holiday: No),(Day name: Monday),(Day period: Night),(Profile: Offline),(Location: Home)}
Top 3 Recommendation Results for user *#152*	
Ground truth	*Multimedia*
PCR-LDA	Multimedia (√), Web, Game
PCR-i	Multimedia (√), Business, Management
CPR	Web, System, Business
Top 3 Recommendation Results for user *#343*	
Ground truth	*Web*
PCR-LDA	Web (√), Multimedia, SNS
PCR-i	Multimedia, Game, Web
CPR	Web (√), System, Business

6 Concluding Remarks

In this paper, we investigated how to exploit user context logs for personalized context-aware recommendation by mining CCPs through topic models. To be specific, first we extract ACP-Feature bags for each user from their historical context logs. Then, we propose to mine users' CCPs through topic models. Finally, we make recommendation according to the given context and the CCP distribution of the given user. The experimental results from a real-world data set clearly show that our proposed recommendation approach can achieve good performance for personalized context-aware recommendation.

References

1. Azzopardi, L., Girolami, M., Risjbergen, K.V.: Investigating the relationship between language model perplexity and ir precision-recall measures. In: SIGIR 2003, pp. 369–370 (2003)
2. Bader, R., Neufeld, E., Woerndl, W., Prinz, V.: Context-aware poi recommendations in an automotive scenario using multi-criteria decision making methods. In: CaRR 2011, pp. 23–30 (2011)
3. Bao, T., Cao, H., Chen, E., Tian, J., Xiong, H.: An unsupervised approach to modeling personalized contexts of mobile users. In: ICDM 2010, pp. 38–47 (2010)
4. Blei, D.M., Ng, A.Y., Jordan, M.I.: Lantent dirichlet allocation. Journal of Machine Learning Research, 993–1022 (2003)
5. Eagle, N., Clauset, A., Quinn, J.A.: Location segmentation, inference and prediction for anticipatory computing. In: AAAI Spring Symposium on Technosocial Predictive Analytics (2009)
6. Griffiths, T.L., Steyvers, M.: Finding scientific topics. Proceedings of National Academy of Science of the USA, 5228–5235 (2004)
7. Heinrich, G.: Paramter stimaion for text analysis. Technical report, University of Lipzig (2009)

8. Hofmann, T.: Probabilistic latent semantic indexing. In: SIGIR 1999, pp. 50–57 (1999)
9. Jae Kim, K., Ahn, H., Jeong, S.: Context-aware recommender systems using data mining techniques. Journal of World Academy of Science, Engineering and Technology 64, 357–362 (2010)
10. Karatzoglou, A., Amatriain, X., Baltrunas, L., Oliver, N.: Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In: RecSys 2010, pp. 79–86 (2010)
11. Nigam, K., McCallum, A.K., Thrun, S., Mitchell, T.: Text classification from labeled and unlabeled documents using EM. Machine Learning 39, 103–134 (2000)
12. Park, M.-H., Hong, J.-H., Cho, S.-B.: Location-Based Recommendation System Using Bayesian User's Preference Model in Mobile Devices. In: Indulska, J., Ma, J., Yang, L.T., Ungerer, T., Cao, J. (eds.) UIC 2007. LNCS, vol. 4611, pp. 1130–1139. Springer, Heidelberg (2007)
13. Sohn, T., Li, K.A., Lee, G., Smith, I., Scott, J., Griswold, W.G.: Place-Its: A Study of Location-Based Reminders on Mobile Phones. In: Beigl, M., Intille, S.S., Rekimoto, J., Tokuda, H. (eds.) UbiComp 2005. LNCS, vol. 3660, pp. 232–250. Springer, Heidelberg (2005)
14. Tung, H.-W., Soo, V.-W.: A personalized restaurant recommender agent for mobile e-service. In: EEE 2004, pp. 259–262 (2004)
15. van Setten, M., Pokraev, S., Koolwaaij, J.: Context-Aware Recommendations in the Mobile Tourist Application COMPASS. In: De Bra, P.M.E., Nejdl, W. (eds.) AH 2004. LNCS, vol. 3137, pp. 235–244. Springer, Heidelberg (2004)
16. Woerndl, W., Schueller, C., Wojtech, R.: A hybrid recommender system for context-aware recommendations of mobile applications. In: ICDE 2007, pp. 871–878 (2007)
17. Zheng, V.W., Cao, B., Zheng, Y., Xie, X., Yang, Q.: Collaborative filtering meets mobile recommendation: A user-centered approach. In: AAAI 2010, pp. 236–241 (2010)