Mining Mobile User Preferences for Personalized Context-Aware Recommendation

HENGSHU ZHU and ENHONG CHEN, University of Science and Technology of China
HUI XIONG, Rutgers University
KUIFEI YU and HUANHUAN CAO, Nokia Research Center
JILEI TIAN, Nokia

Recent advances in mobile devices and their sensing capabilities have enabled the collection of rich contextual information and mobile device usage records through the device logs. These context-rich logs open a venue for mining the personal preferences of mobile users under varying contexts and thus enabling the development of personalized context-aware recommendation and other related services, such as mobile online advertising. In this article, we illustrate how to extract personal context-aware preferences from the context-rich device logs, or ***context logs*** for short, and exploit these identified preferences for building personalized context-aware recommender systems. A critical challenge along this line is that the context log of each individual user may not contain sufficient data for mining his or her context-aware preferences. Therefore, we propose to first learn common context-aware preferences from the context logs of many users. Then, the preference of each user can be represented as a distribution of these common context-aware preferences. Specifically, we develop two approaches for mining common context-aware preferences based on two different assumptions, namely, context-independent and context-dependent assumptions, which can fit into different application scenarios. Finally, extensive experiments on a real-world dataset show that both approaches are effective and outperform baselines with respect to mining personal context-aware preferences for mobile users.

Categories and Subject Descriptors: H.2.8.d [**Information Technology and Systems**]: Database Applications—*Data mining*; I.2.6 [**Artificial Intelligence**]: Learning

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Context aware, personalized recommendation, mobile user

ACM Reference Format:
Hengshu Zhu, Enhong Chen, Hui Xiong, Kuifei Yu, Huanhuan Cao, and Jilei Tian. 2014. Mining mobile user preferences for personalized context-aware recommendation. ACM Trans. Intell. Syst. Technol. 5, 4, Article 58 (December 2014), 27 pages.
DOI: http://dx.doi.org/10.1145/2532515

1. INTRODUCTION

With the rapid development of the mobile industry, the mobile platform, such as smartphones and tablets, has become one of the most important media for social

This is a substantially extended and revised version of Zhu et al. [2012b], which appears in Proceedings of the 12th IEEE International Conference on Data Mining (ICDM'2012).
Authors' addresses: H. Zhu and E. Chen are with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui, 230027, China; emails: zhs@mail.ustc.edu.cn, cheneh@ustc.edu.cn; H. Xiong is with the Management Science and Information Systems Department, Rutgers Business School, Rutgers University, Newark, NJ 07102 USA; email: hxiong@rutgers.edu; K. Yu and H. Cao are with the Nokia Research Center, Beijing, 100010, China; email: yukuifei@gmail.com, happia.cao@gmail.com; J. Tian is with the Nokia, Beijing, 100010, China; email: jilei.tian@nokia.com.

DOI: http://dx.doi.org/10.1145/2532515

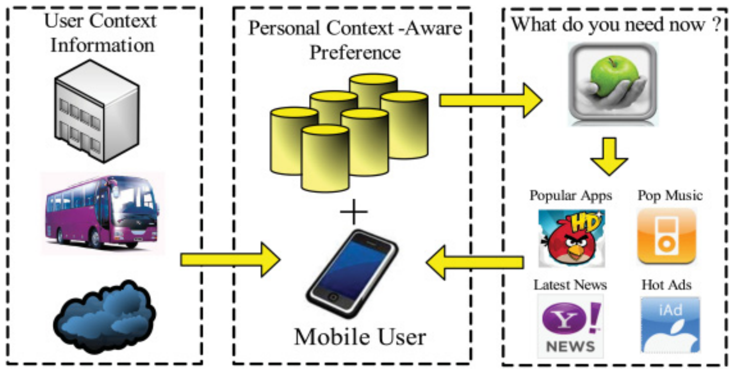

Fig. 1. The personalized context-aware recommendation services for mobile users.

entertainment and information acquisition. Meanwhile, the advances in mobile devices enable them to be equipped with a rich set of context sensors, such as GPS sensors, 3D accelerometers, and optical sensors. These context sensors can capture the rich contextual information of mobile users, and thus enable a wide range of context-aware services, such as context-aware tour guide [Emrich et al. 2009], location-based reminder [Sohn et al. 2005], and context-aware recommendation [Bader et al. 2011; Liu et al. 2013; jae Kim et al. 2010; Karatzoglou et al. 2010; Woerndl et al. 2007]. In fact, the contextual information and corresponding usage records (e.g., browsing websites and playing games) can be recorded in context-rich device logs, or ***context logs*** for short, which can be used for mining the personal context-aware preferences of mobile users, that is, which category of contents is preferred by a particular user under a certain context. Particularly, mining such preferences is a fundamental work for understanding the behaviors of mobile users, and thus enables better context-aware services. Specifically, by considering both the personal context-aware preferences and the current contexts of users, it is possible to provide personalized context-aware recommendation and other related services for mobile users, such as mobile online advertising. Figure 1 shows some examples of how to exploit contextual information for recommendation services. Indeed, the personalized context-aware recommendation services can provide better user experiences than traditional context-aware recommender systems, which only consider the contextual information but not different users' preferences under the same context [Zhu et al. 2012b]. For instance, the following two examples intuitively illustrate how the context-aware recommendation services can improve the user experiences.

Example 1.1 (*Context-Aware Content Pushing*). Content pushing is an important function of smart mobile operating systems, which aims to automatically deliver the right content to mobile users and is widely used for content recommendation and intelligent advertising [Podnar et al. 2002]. Particularly, suppose that a mobile user Kate would like to play her mobile phone while taking the bus. Through the analysis of the sensing information collected on her smartphone, a context-aware recommender system could discover that Kate is taking a bus (sensed by 3D accelerometers [Kwapisz et al. 2011]) from her workplace to home (sensed by GPS or cell ID combined with data mining on historical trajectories of users [Eagle et al. 2009; Zheng et al. 2009]) on a Monday evening (sensed by the system clock). Therefore, with respect to Kate's personal context-aware preference, for example, that Kate often listens to R&B music under the same context, the recommender system may push some new R&B music and related advertisements to Kate.

Example 1.2 (*Context-Aware Mobile Advertising*). Mobile advertising, such as banner or in-app advertising, has become one of the most important components of online advertising campaigns. In fact, a crucial need of mobile advertisers is to understand the behavior and preferences of mobile users. Therefore, with the help of the personalized context-aware recommender system, the advertiser can better design and deliver advertisements to mobile users. For example, through mining users' context logs, the recommender system may discover that a mobile user, Tom, has two typical context-aware preferences "playing games every evening at home" and "using restaurant review apps every workday afternoon near the office"; thus, it can suggest advertisers to deliver advertisements of popular games and restaurants to Tom when he is using his mobile phone in corresponding contexts. Furthermore, such context-aware preferences can also be used for reranking the sponsored advertisements with respect to users' different contexts.

In recent years, although many researchers studied the problem of personalized context-aware recommendation [Zheng et al. 2010; Park et al. 2007; Ge et al. 2011; Liu et al. 2011; jae Kim et al. 2010] and proposed some approaches for mining personal context-aware preferences, most of them did not take into account context-rich information in their approaches. Also, some of these studies are based on item ratings generated by users under different contexts, which are difficult to obtain in practice. In contrast, usage records in context-rich device logs are a rich resource for mining personalized context-aware user preferences. However, how to mine context-aware preferences from context-rich logs for developing context-aware recommender systems is still underaddressed.

To this end, in this article, we propose a novel approach for mining personal context-aware preferences from context-rich device logs of mobile users. A critical challenge for mining personal context-aware preferences is that the context log of each individual user usually does not contain sufficient training information. As a result, it can be difficult to learn personal context-aware preferences if we only use the context log of each individual user. Therefore, we propose to first find common context-aware preferences from the context logs of many users and then represent the context-aware preference of each user as a distribution of common context-aware preferences. Moreover, on the basis of two different assumptions about context data dependency, we propose two methods for mining common context-aware preferences. The first one is more efficient but sacrifices a little performance, while the second one needs more training time but has better performances. Figure 2 illustrates the overview of the proposed approach and how the mined preferences are used for predicting the preferred categories of contents for a given mobile user under a certain context. Specifically, the contributions of this article are summarized as follows.

First, we propose a novel approach for mining the personal context-aware preferences for mobile users through the analysis of context-rich device logs. Specifically, we propose to first mine common context-aware preferences from the context logs of many users and then represent the personal context-aware preference of each user as a distribution of common context-aware preferences. The mined personal context-aware preferences can enable the development of personalized context-aware recommender systems and other related services, such as mobile online advertising.

Second, we design three effective methods for mining common context-aware preferences based on two different assumptions about context data dependency. If context data are assumed to be conditionally independent, we propose to mine common context-aware preferences through topic models. Otherwise, if context data are assumed to be dependent, we propose to exploit the constraint-based Matrix Factorization techniques

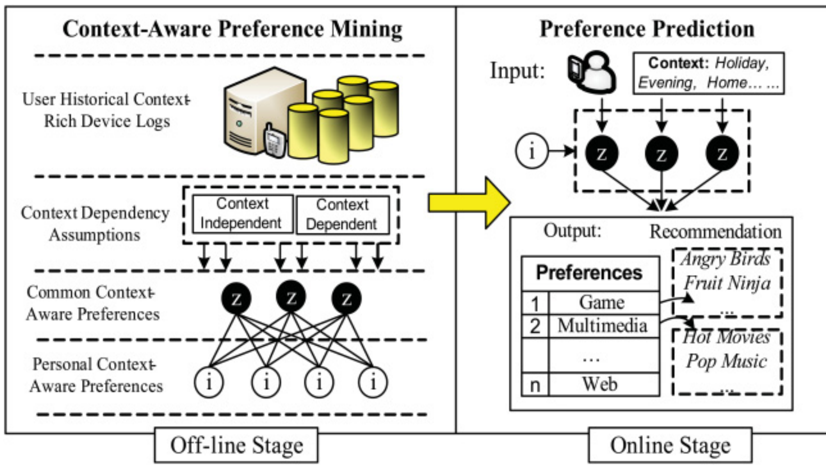

Fig. 2. The overview of the proposed approach for mining personal context-aware preferences and how the mined preferences are used for predicting the preferred categories of contents for a given mobile user under a certain context.

for mining common context-aware preferences and only consider those contexts that are relevant to content usage for reducing the computation complexity.

Finally, we evaluate the proposed approach using a real-world dataset with context logs collected from 443 mobile phone users. In total, there are more than 8.8 million context records. This dataset contains much more context-rich information and is much bigger than those reported in previous works on context log mining [Bao et al. 2010; Cao et al. 2010]. The experimental results clearly demonstrate the effectiveness of the proposed approach and indicate some inspiring findings.

Overview. The remainder of this article is organized as follows. In Section 2, we introduce the details about context logs and give the overview of the proposed approach. Section 3 and Section 4 present our novel approaches for mining common context-aware preferences of many users on the basis of two different assumptions, respectively. In Section 5, we report the experimental results on a real-world dataset. Section 6 provides a brief review of related works. Finally, in Section 7, we conclude the article.

2. MINING PERSONAL CONTEXT-AWARE PREFERENCES FROM CONTEXT LOGS

Smart devices can capture the historical context data and the corresponding usage records of users through multiple sensors and record them in context logs. For example, Table I shows a toy context log, which contains several ***context records***, and each context record consists of a timestamp, the most detailed context at that time, and the corresponding usage record. 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***. Moreover, usage records can be empty (denoted as "Null") because a user does not always use the mobile device.

Note that, in Table I, raw locations in context data, such as GPS coordinates or cell IDs, have been transformed into semantic locations such as "Home" and "Workplace" by some location mining approaches (e.g., [Eagle et al. 2009]). The basic idea of these approaches is to find the clusters of user locations and recognize their semantic meaning by a time pattern analysis. Moreover, we also map the raw usage records to the usage records of particular categories of contents via some mobile app classification approaches (e.g., [Zhu et al. 2012a]). For example, we can map two raw usage records

Table I. A Toy Context-Rich Device Log Collected by Mobile Devices

Timestamp	Context	Activity Records
t_1	{(*Day name: Monday*), (*Time range: AM8:00-9:00*)), (*Profile: General*), (*Battery: 5*), (*Location: Home*)}	Null
t_2	{(*Day name: Monday*),(*Time range: AM8:00-9:00*)), (*Profile: General*),(*Battery: 5*),(*Location: On the way*)}	Play action games (*Fruit Ninja*)
t_3	{(*Day name: Monday*),(*Time range: AM8:00-9:00*)), (*Profile: General*),(*Battery: 5*),(*Location: On the way*)}	Null
......	......	
t_{359}	{(*Day name: Monday*),(*Time range: AM10:00-11:00*), (*Profile: Meeting*),(*Battery: 4*),(*Location: Work place*)}	Null
t_{360}	{(*Day name: Monday*),(*Time range: AM10:00-11:00*), (*Profile: Meeting*),(*Battery: 4*),(*Location: Work place*)}	Browsing sports websites (*www.nba.com*)
......	......	
t_{448}	{(*Day name: Monday*),(*Time range: AM11:00-12:00*), (*Profile: General*),(*Battery: 4*),(*Location: Work place*)}	Play with SNS (*Facebook*)
t_{449}	{(*Day name: Monday*),(*Time range: AM11:00-12:00*), (*Profile: General*),(*Battery: 4*),(*Location: Work place*)}	Null

"*Play Angry Birds*" and "*Play Fruit Ninja*" to the usage records of content category "*Action Games.*"

In this way, the context data and usage records in context logs are normalized and the data sparseness problem is somewhat alleviated. This helps the task of personal context-aware preference mining.

Intuitively, context logs contain rich information about content usage given particular contexts and can be used for mining the personal context-aware preferences of users. However, the context log of each individual user is usually too sparse for this task. This is also demonstrated by the experiments on a real-world dataset in the experimental section. The main reason is that, while the context logs of individual users may contain many context records, only a small proportion of them have nonempty usage records that can be used as a meaningful mining source. To that end, we propose a novel approach for mining personal context-aware preferences as follows.

The basic idea is first mining common context-aware preferences from the context logs of many users and then representing each user's context-aware preference by a distribution of common context-aware preferences. Let us denote the variable of common context-aware preference as z, and the conditional probability that a user u prefers the content category c given a context C can be represented as

$$P(c|C,u) = \frac{P(c,C|u)\cdot P(u)}{P(C,u)} \propto P(c,C|u)$$
$$\propto \sum_z P(c,C,z|u) \propto \sum_z (P(c,C|z)\cdot P(z|u)),$$

where we assume that a user's preference given a context only relies on the common context-aware preferences followed by many users, that is, $P(c, C|z)$, and his or her personal context-aware preference expressed by a distribution of common context-aware preferences, that is, $P(z|u)$. Then the task is converted to learn $P(c, C|z)$ and $P(z|u)$ from many users' context logs. Specifically, given a user u and context C, both $P(u)$ and $P(C, u)$ are constant, and thus we have $P(c|C, u) \propto P(c, C|u)$ in the previous equation.

After mining the personal context-aware preference of each mobile user, we predict which category of contents will be preferred for a given user according to the corresponding context. Specially, we first rank content categories according to the probability

$P(c|C,u)$ of each content category c, and then we infer user-preferred content category c^* by $c^* = \arg\max_c P(c|C,u)$ and recommend corresponding contents. For example, if we infer the user would like "*Action Games,*" we will recommend some popular action games to the user.

We observe that modeling and mining common context-aware preferences rely on the assumption about context data dependency. Basically, we can have two different assumptions about context data dependency as follows.

The first assumption is that different types of context data are conditionally independent given a particular common context-aware preference, which is relatively strong but can simplify the problem. For example, under such an assumption, given a context "{(*Time range: PM10:00-11:00*), (*Location: Home*)}" and a mobile user u, if we can infer the latent common context-aware preference distribution of u, we only need to consider which category of contents u may prefer under the context (*Time range: PM10:00-11:00*) and the context (*Location: Home*) given each common context-aware preference, but do not need to consider which category of contents u may prefer given the co-occurrence of (*Time range: PM10:00-11:00*) and (*Location: Home*) given each common context-aware preference.

The second assumption is that different types of context data are mutually dependent, which is relatively weak and may be more proper in practice. However, such an assumption makes it more difficult for modeling context-aware preferences. For example, under such an assumption, given the previously mentioned context, we have to consider the co-occurrence of (*Time range: PM10:00-11:00*) and (*Location: Home*) when making a preference prediction. Obviously, the corresponding models may be more complex than the ones based on the first assumption. In this article, we propose different approaches based on the previous two assumptions and conduct extensive experiments to evaluate them. The details of our novel approaches based on two different assumptions are presented in the following two sections, respectively.

3. CONTEXT-AWARE PREFERENCE MINING BASED ON CONTEXT CONDITIONAL INDEPENDENCY ASSUMPTION

We first propose a method based on the assumption that different types of context data are conditional independent given a particular common context-aware preference. Under such an assumption, given a context $C = \{p_1, p_2, \ldots, p_l\}$, where p_i denotes an *atomic context*, that is, a contextual feature–value pair, the probability that a user u prefers content category c can be represented as

$$P(c|C,u) \propto \sum_z (P(c,C|z)\cdot P(z|u)) \propto \sum_z \left(\prod_{p_i \in C} P(c,p_i|z)\cdot P(z|u)\right).$$

Therefore, the problem is further converted to learn $P(c, p_i|z)$ and $P(z|u)$ from many users' context logs, which can be solved by widely used topic models. In this section, we present how to utilize topic models for mining common context-aware preferences by estimating $P(c, p_i|z)$ and $P(z|u)$. For simplicity, we refer to the co-occurrence of a usage of a content in category c and the corresponding contextual feature–value pair p_i, that is, (c, p_i), as an ***Atomic Context-aware Preference feature***, or ***ACP-feature*** for short. For example, $(IsHoliday?: Yes) \rightarrow$ *Games* is an ACP-feature, where $p = (IsHoliday?: Yes)$ and $c =$ *Games*.

3.1. 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 $\{w_{i,j}\}$ that are generated by these topics. Intuitively, if we take

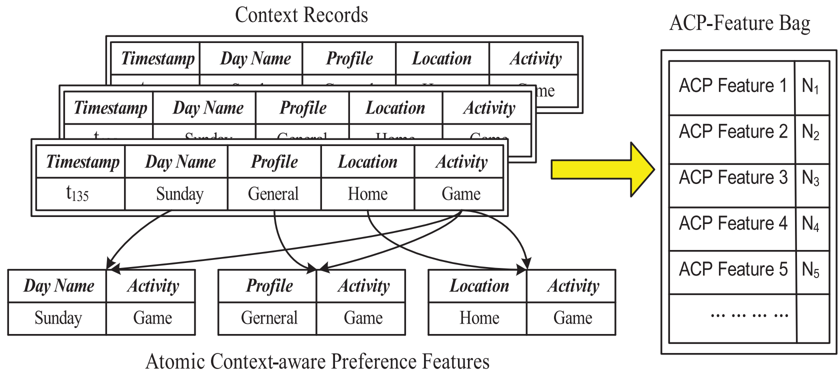

Fig. 3. The generation process of ACP-feature bag from user's context records.

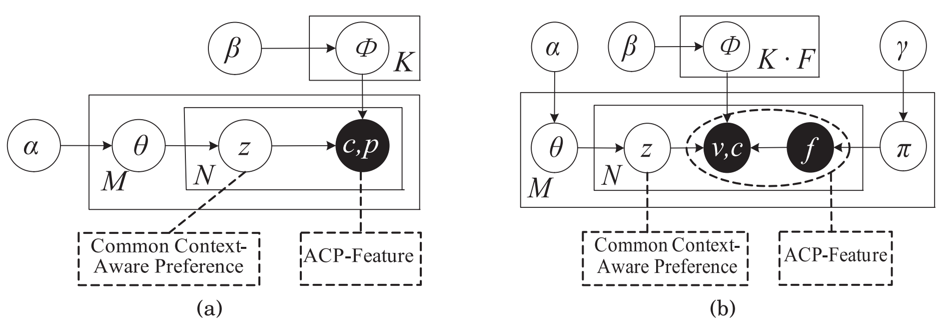

Fig. 4. The graphic representation of modeling ACP-feature bags by (a) LDA topic model and (b) LDAC topic model.

ACP-features as words, take context logs as bags of ACP-features to correspond documents, and take common context-aware preferences as topics, we can take advantage of topic models to learn common context-aware preferences from many users' context logs.

Since raw context logs are not naturally in the form of bags of ACP-features, we need to extract bags of ACP-features from them as training data. Especially, we first remove all context records without any usage record and then extract the ACP-feature from the remaining ones. Given a context record $<Tid, C, c>$, where Tid denotes a timestamp, $C = \{p_1, p_2, \ldots, p_l\}$ denotes a context, and c denotes the category of the used content in the usage record, we can extract l ACP-features, namely, (c, p_1),(c, p_2), $\ldots$, (c, p_l). For simplicity, we refer to the bag of ACP-features extracted from user u's context log as the ***ACP-feature bag*** of u. Figure 3 shows the example of the generation process of the ACP-feature bag from context records.

3.1.1. LDA-Based Context Modeling. Among several existing topic models, in this article, we first leverage the widely used Latent Dirichlet Allocation model (LDA) [Blei et al. 2003]. According to LDA, 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. Second, a prior common context-aware preference distribution θ_i is generated from a prior Dirichlet distribution α for each user u_i. Then, for generating the jth ACP-feature in d_i denoted as $w_{i,j}$, the model first generates a common context-aware preference z from θ_i and then generates $w_{i,j}$ from ϕ_z. Figure 4(a) shows the graphic representation of modeling ACP-feature bags by the LDA model.

The process of LDA model training is to learn the proper latent variables θ and ϕ to maximize the posterior distribution of the observed ACP-feature bags, that is, $P(u|\alpha, \beta, \theta, \phi)$. In this article, we take advantage of a Markov chain Monte Carlo method named Gibbs sampling [Griffiths and Steyvers 2004] for training LDA model. This method begins with a random assignment of common context-aware preferences 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 common context-aware preference to each ACP-feature, which is conditional on the assignment of all other ACP-features. Then a new assignment of common context-aware preferences 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 common context-aware preference and we can obtain the estimation of $P(z|u)$ and $P(c, p|z)$ as follows:

$$P(z|u) = P(z|d_u) = \frac{n_{u,z} + \alpha_z}{\sum_i^K n_{u,z_i} + \sum_i^K \alpha_{z_i}}, \tag{1}$$

$$P(c, p|z) = \frac{n_{z,c,p} + \beta_{c,p}}{\sum_i^M n_{z,(c,p)_i} + \sum_i^M \beta_{(c,p)_i}}, \tag{2}$$

where n_{u,z_i} is the frequency of ACP-features in u that have been assigned to common context-aware preference z_i, and $n_{z,(c,p)_i}$ is the frequency of the ith ACP-feature that has been assigned to common context-aware preference z.

3.1.2. LDAC-Based Context Modeling. Although LDA can model the ACP-feature bags in an intuitive way, from some real-world observations we find that the generation of ACP-features are decided not only by latent common context-aware preferences but also by their internal contextual features. For example, Bluetooth information can only be obtained when the user opens the Bluetooth sensor, and location information often cannot be obtained in underground subways due to the lack of GPS/cell ID information. Therefore, to more accurately model context information, in this article, we also leverage the extended Latent Dirichlet Allocation on Context Model (LDAC) [Bao et al. 2010] for mining latent common context-aware preferences.

In the LDAC model, the ACP-feature bags of user u_i denoted as d_i are generated as follows. First, a prior common context-aware preference distribution θ_{d_i} is generated from a prior Dirichlet distribution α. Second, a prior contextual feature distribution π_{d_i} is generated from a prior Dirichlet distribution γ. Then, for the jth ACP-feature in d_i, a common context-aware preference $z_{d_i,j}$ is generated from θ_{d_i}, a contextual feature $f_{d_i,j}$ is generated from π_{d_i}, and the content category with contextual feature value of $f_{d_i,j}$ denoted as $(c, v_{d_i,j})$ is generated from the distribution $\phi_{z_{d_i,j}, f_{d_i,j}}$. Moreover, there are totally $K \times F$ prior distributions of ACP-features $\{\phi_{k,f}\}$ that follow a Dirichlet distribution β, where F is the number of unique contextual features. Figure 4(b) shows the graphical representation of the LDAC model. The process of LDAC model training is to learn the proper latent variables to maximize the posterior distribution of the observed ACP-feature bags, that is, $P(d, \theta, z_d, \pi_d, \Phi|\alpha, \beta, \gamma)$.

In this article, we leverage the Gibbs sampling-based approach introduced in Bao et al. [2010] to train the LDAC model. After the training process, we can obtain the probabilities $P(z|u)$ by Equation (1) and $P(c, p|z)$ as follows: [1]

$$P(c, p|z_k) = P(c, v_p| f_p, z)P(f_p), \tag{3}$$

[1]The detailed inference can be found in the appendix.

where $p = (f_p : v_p)$ and

$$P(c, v_p | f_p, z) = \frac{n_{z,c,f_p,v_p} + \beta_{c,v_p}}{\sum_v n_{z,c,f_p,v} + \sum_{v \in V_{f_m}} \beta_v},$$

$$P(f_p) = \frac{\sum_z \sum_v n_{k,c,f_p,v} + \gamma_{f_m}}{\sum_f \sum_z \sum_v n_{z,c,f_p,v} + \sum_f \gamma_f},$$

where n_{z,c,f_p,v_p} is the frequency of ACP-feature (c, p) that has been assigned to common context-aware preference z, and V_{f_m} is the number of values for contextual feature f_m.

3.2. Selecting the Number of Common Context-Aware Preferences

Both LDA and LDAC models need a predefined parameter K to determine the number of common context-aware preferences. In this article, we utilize the method proposed in Bao et al. [2010] to estimate K. To be specific, we first empirically define a topic number range $[K_{min}, K_{max}]$ and then select two groups of ACP-feature bags as the training set S_1 and the test set S_2, respectively. Then we can determine K by the corresponding perplexity [Azzopardi et al. 2003; Blei et al. 2003] of the test set S_2. Here we take the LDA model as an example, and the perplexity is defined as follows:

$$Perplexity(S_2) = Exp\left\{-\frac{\sum_{d_i \in S_2} log\{P(d_i|S_1)\}}{\sum_{d_i \in S_2} N_{d_i}}\right\},$$

where N_{d_i} indicates the number of ACP-features in d_i and

$$P(d_i|S_1) = \prod_{(c,p)\in d_i} P(c, p|S_1) = \prod_{(c,p)\in d_i} \sum_{j=1}^{K} (P(c, p|z_j)P(z_j|S_1)).$$

Herein, $P(c, p|z_j)$ can be obtained after model training, and $P(z_j|S_1)$ can be estimated by $\frac{n_{d_i,j}+\alpha}{\sum_{k=1}^{K} n_{d_i,k}+\alpha}$, where $n_{d_i,j}$ indicates the number of ACP-features labeled with z_j in d_i.

The perplexity is widely used to index the modeling performance. The smaller the perplexity is, the better modeling performance it implies. However, the perplexity may consistently drop with the increase of K in practical datasets. To avoid the overfitting problem, we cannot only utilize the minimum perplexity as the metric for determining K [Azzopardi et al. 2003; Blei et al. 2003]. A complementary method is to define a decline rate ζ of perplexity and stop seeking better K if the decline rate of perplexity is less than ζ. In our experiments, we set ζ to be 10% according to Bao et al. [2010].

4. CONTEXT-AWARE PREFERENCE MINING BASED ON CONTEXT DEPENDENCY ASSUMPTION

Since it may be relatively strong to assume that different types of context data are conditionally independent, we also propose a method for mining common context-aware preferences based on the assumption that different types of context data are mutually dependent. Under such an assumption, we cannot decompose contexts into atomic contexts and need to learn $P(c, C|z)$ directly from user context logs. A major challenge is that we cannot learn all conditional distributions $P(c, C|z)$ for all C simply because the number of unique C is exponential to the number of unique contextual feature–value pairs and we will suffer the assemble explosion problem if we learn all of them. Fortunately, we observe that usually not all parts of a context are relevant to content usage and thus the corresponding preferences of content categories. For example, given a context "{(*Day name: Sunday*), (*Day period: Evening*), (*Location: Home*), (*Battery* *level: Low*)},” we may not be able to demonstrate that the whole context is relevant to a particular content category through the analysis of context logs. Instead, we may find that context logs show that some parts of them, such as “{(*Day name: Sunday*), (*Day period: Evening*)” and “{(*Day period: Evening*), (*Location: Home*)},” are indeed relevant to a particular content category. To this end, an intuitive idea is to only consider the content-relevant parts of contexts for predicting personalized context-aware preferences of content categories. These content-relevant parts can be referred to as ***content-relevant contexts*** for simplicity. Along this line, we only need to learn the conditional distributions $P(c, C^r|z)$ and $P(z|u)$, where C^r denotes a content-relevant context. Moreover, given a context C, it can be divided into two cases to calculate $P(c, C|z)$.

First, if C contains some content-relevant contexts, we can calculate $P(c, C|z)$ directly by its maximal subcontexts, which are also content-relevant contexts, as follows:

$$P(c, C|z) = \frac{1}{|C^r_{max}|} \times \sum_{C^r_{max}} P\left(c, C^r_{max}|z\right),$$

where C^r_{max} denotes a maximal content-relevant subcontext contained by C, and $|C^r_{max}|$ indicates of the number of C^r_{max}. Take the previous context C={(*Day name: Sunday*), (*Day period: Evening*), (*Location: Home*), (*Battery level: Low*)} for example; suppose that there are two maximum content-relevant subcontexts C^r_1={(*Day name: Sunday*), (*Day period: Evening*)} and C^r_2={(*Day period: Evening*), (*Location: Home*)} in it: we will recommend corresponding contents to the given user by considering his or her preference under both C^r_1 and C^r_2 equally.

Second, if C does not contain any content-relevant context, we can estimate $P(c, C|z)$ by normalizing the probabilistic space of the joint distribution of c and C conditional on z. Especially, we let

$$P(c, C^{\phi}|z) = \frac{1}{N_{\phi} \cdot N_c} \times \left(1 - \sum_{a} \sum_{C^{\triangle}} P(c, C^{\triangle}|z)\right),$$

where C^{ϕ} denotes a context without any content-relevant subcontext, N_{ϕ} indicates the total number of C^{ϕ}, N_c indicates the total number of unique content categories, and $C^{\triangle}$ denotes a context with at least one content-relevant subcontext. Actually, we do not need to calculate $P(c, C|z)$ in this case because it is the same with varying c and cannot help to make a recommendation decision. In practice, we do not recommend any content in this case.

Therefore, the original problem is divided into two subproblems, namely, *how to discover those content-relevant contexts* and *how to learn common context-aware preferences and user personal distributions of common context-aware preferences, that is,* $P(c, C^r|z)$ *and* $P(z|u)$. The solutions for the two subproblems are presented in the following sections in detail, respectively.

4.1. Discovering the Content-Relevant Context

An intuitive way of discovering the context relevant to some content categories is by mining association rules [Agrawal and Srikant 1994] between them with predefined minimum supports and minimum confidences. Therefore, given a content-relevant context C^r and a content category c, we have $P(c, C^r|u) = P(c|C^r, u)P(C^r|u)$, where $P(c|C^r, u)$ can be estimated by the corresponding confidence of the association “$C^r \longrightarrow c$” and $P(C^r|u)$ can be estimated by $\frac{Support(C^r)}{N_r}$, where N_r indicates the total number of context records in the context log of user u.

Table II. Examples of Association Rules Mined from Context Logs

(*Is holiday? Yes*) (*Time Period: Evening*) (*Location: Home*) ⇒ **Game**
(*Is holiday? Yes*) (*Time Period: Evening*) (*Charging State: Charging*) ⇒ **Game**
(*Time Period: Morning*) (*Location: Workplace*)⇒ **Business**
(*Time Period: Evening*) (*Location: Moving*) (*Profile: Silent*) ⇒ **Multimedia**
(*Time Period: Evening*) (*Location: Home*) ⇒ **Web**

However, as pointed out by Cao et al. [2010], the amounts of context data and user activity records are usually extremely unbalanced, which makes it difficult to mine such association rules through traditional association rule mining approaches. An alternative approach is only leveraging the context records with nonempty activity records. However, it will lose the discriminative information on how likely no activity will be done under a particular context. Fortunately, some researchers have studied this problem and proposed some novel algorithms for mining such association rules. For example, Cao et al. [2010] proposed a novel algorithm called GCPM (Generating Candidates for behavior Pattern Mining) for mining such association rules, which are referred to as behavior patterns in their work, by utilizing different ways of calculating supports and confidences. In an incremental work of Cao et al. [2010], Li et al. [2012] proposed a more efficient algorithm named BP-Growth for this problem. In this article, we leverage the BP-Growth algorithm for mining such association rules. The basic idea of the algorithm is partitioning the original context logs into smaller subcontext logs for reducing the mining space and mining frequent association rules in these subcontext logs. Table II illustrates some examples of association rules mined from context logs. In fact, these association rules clearly demonstrate the personal context-aware preferences of mobile users. For example, the context "(*Time Period: Evening*) (*Location: Home*)" may imply a leisure time at home, and thus the user has a preference "Web" under such context.

It is worth noting that the mining is performed on individual users' context logs because merging all context logs may normalize the associations between contexts and content categories. For example, given that several users usually play action games on the bus and several other users usually play other games on the bus, if we try to mine the associations between contexts and content category by merging all users' context logs, we may falsely conclude that "*On the bus*" has no significant relevance with any content category. In contrast, we can discover that "*On the bus*" is relevant to both action games and other games according to different people by taking into account each user's context log separately.

4.2. Mining Common Context-Aware Preferences Through Constraint-Based Bayesian Matrix Factorization

After finding content-relevant contexts, the remaining task is to learn common context-aware preferences and user personal distributions of common context-aware preferences, that is, $P(c, C^r|z)$ and $P(z|u)$. By building a matrix of $P(c, C^r|u)$, where each column denotes a probabilistic distribution of different (c, C^r) pairs for a given user u, we can convert this task into a matrix factorization problem as follows:

$$\mathbf{\Omega}_{N\times M} = \mathbf{\Phi}_{N\times K} \times \mathbf{\Theta}_{K\times M} + \mathbf{N}_{N\times M},$$

where N indicates the number of unique (c, C^r) pairs, M indicates the number of users, and K indicates the number of common context-aware preferences. To be specific, $\mathbf{\Omega}$ denotes the observed matrix of $P(c, C^r|u)$, $\phi_{ik} \in \mathbf{\Phi}(1 \le i \le N, 1 \le k \le K)$ denotes the probability $P(c, C^r|z_k)$, $\theta_{kj} \in \mathbf{\Theta}(1 \le k \le K, 1 \le j \le M)$ denotes the probability

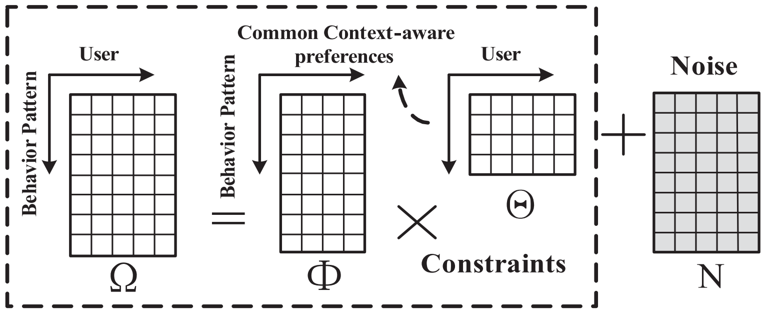

Fig. 5. Mining common context-aware preferences by constrained matrix factorization.

$P(z_k|u_j)$, and the matrix $\mathbf{N}$ denotes the residual noise information. Moreover, the matrix factorization task has two additional constraints for possible solutions as follows: (1) all elements in matrix $\mathbf{\Phi}$ and $\mathbf{\Theta}$ should be nonnegative values, and (2) $\forall_{j:1\leq j\leq M} \sum_{k=1}^{K} \theta_{kj} = 1$ and $\forall_{k:1\leq k\leq K} \sum_{i=1}^{N} \phi_{ik} = 1$, which are both obvious since each column of $\mathbf{\Phi}$ and $\mathbf{\Theta}$ denotes a probabilistic distribution. Figure 5 shows the process of mining common context-aware preferences through constrained matrix factorization.

According to the earlier problem statement and constraints, the objective of our matrix factorization task is to find a possible solution for matrix $\mathbf{\Phi}$, $\mathbf{\Theta}$, and $\mathbf{N}$. In this article, we propose to leverage a constraint-based Bayesian Matrix Factorization model [Schmidt 2009] for resolving this problem. In this model, we can perform matrix factorization with multiple inequality and equality constraints. Specifically, we aim to infer the posterior probabilistic distributions of $\mathbf{\Phi}$ and $\mathbf{\Theta}$ under a set of model assumptions, which are specified by the likelihood function $P(\mathbf{\Omega}|\mathbf{\Phi}, \mathbf{\Theta}, \mathbf{N})$. The likelihood function denotes the probability of the observed data matrix $\mathbf{\Omega}$ given priors $P(\mathbf{\Phi}, \mathbf{\Theta})$ and $P(\mathbf{N})$. According to Schmidt [2009], to perform efficient inference based on Gibbs sampling, we select priors as follows. First, we select an i.i.d. zero mean Gaussian noise model as follows:

$$P(n_{ij}) = N(n_{ij}|0, \nu_{ij}) = \frac{1}{\sqrt{2\pi\nu_{ij}}} exp\left(-\frac{n_{ij}^2}{2\nu_{ij}}\right), \tag{4}$$

where parameter ν_{ij} satisfies a conjugate inverse-gamma prior that satisfies

$$P(\nu_{ij}) = IG(\nu_{ij}|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \nu_{ij}^{-(\alpha+1)} exp\left(\frac{-\beta}{\nu_{ij}}\right).$$

Then, we select a Gaussian prior over $\mathbf{\Phi}$ and $\mathbf{\Theta}$ subject to inequality constraints $\mathbf{Q}$ and equality constraints $\mathbf{R}$ as

$$P(\overrightarrow{\phi}, \overrightarrow{\theta}) \propto \begin{cases} \mathrm{N}\left(\begin{bmatrix} \overrightarrow{\phi} \\ \overrightarrow{\theta} \end{bmatrix} \middle| \underbrace{\begin{bmatrix} \mu_\phi \\ \mu_\theta \end{bmatrix}}_{\mu}, \underbrace{\begin{bmatrix} \mathbf{\Sigma}_\phi & \mathbf{\Sigma}_{\phi\theta} \\ \mathbf{\Sigma}_{\phi\theta}^{\mathbf{T}} & \mathbf{\Sigma}_s \end{bmatrix}}_{\mathbf{\Sigma}} \right), & if\ \mathbf{Q}(\overrightarrow{\phi}, \overrightarrow{\theta}) \leq 0, \\ & \mathbf{R}(\overrightarrow{\phi}, \overrightarrow{\theta}) = 0, \\ 0, & otherwise, \end{cases} \tag{5}$$

where $\overrightarrow{\phi} = (\phi_{11}, \phi_{12}, \ldots, \phi_{NK})^{\mathbf{T}}$ and $\overrightarrow{\theta} = (\theta_{11}, \theta_{12}, \ldots, \theta_{KM})^{\mathbf{T}}$.

With the previous definitions, we can utilize Gibbs sampling methods to estimate the posterior distributions as follows. In the first round of sampling, we randomly assign values for $\overrightarrow{\phi}$ and $\overrightarrow{\theta}$ according to the two constraints to initialize the state of

Markov chain. Then, we calculate the density of noise variance $P(\nu_{ij}|\overrightarrow{\phi}, \overrightarrow{\theta})$ by inverse-gamma distribution due to the choice of conjugate prior (i.e., Equation (4)). Next, we can estimate $P(\overrightarrow{\phi}|\mathbf{\Omega}, \overrightarrow{\theta}, \mathbf{N})$ and $P(\overrightarrow{\theta}|\mathbf{\Omega}, \overrightarrow{\phi}, \mathbf{N})$ from the constraint Gaussian density (i.e., Equation 5). Finally, we regenerate values for $\overrightarrow{\phi}$ and $\overrightarrow{\theta}$ according to the new posterior probabilities to score a new state of Markov chain. After many rounds of iterations, the results of matrixes $\mathbf{\Phi}$ and $\mathbf{\Theta}$ will converge.

4.3. Selecting the Number of Common Context-Aware Preferences

As mentioned in Section 3, an important problem for mining common context-aware preferences is to select a proper number of common context-aware preferences. For the Matrix Factorization-based approach for mining common context-aware preferences, we utilize the Chib's method introduced in Schmidt et al. [2009] to infer the proper number of common context-aware preferences. To be specific, in the Bayesian framework, model selection can be performed by evaluating the marginal likelihood $P(\mathbf{\Omega})$. The Chib's method is based on the Bayes relation $P(\mathbf{\Omega}) = \frac{P(\mathbf{\Omega}|\mathbf{\Theta})P(\mathbf{\Theta})}{P(\mathbf{\Theta}|\mathbf{\Omega})}$, where the numerator can be easily estimated by the trained model and the key problem is to estimate $P(\mathbf{\Theta}|\mathbf{\Omega})$. Denoting each row of $\mathbf{\Theta}$ as $\mathbf{\Theta}_i$, we can calculate the denominator as $P(\mathbf{\Theta}|\mathbf{\Omega}) = P(\mathbf{\Theta}_1|\mathbf{\Omega}) \times P(\mathbf{\Theta}_2|\mathbf{\Theta}_1, \mathbf{\Omega}) \times \cdots \times P(\mathbf{\Theta}_K|\mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_{K-1}, \mathbf{\Omega})$. After R rounds of Gibbs sampling in our approach, we can estimate each term by averaging over the conditional density $P(\mathbf{\Theta}_K|\mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_{K-1}, \mathbf{\Omega}) \approx \frac{1}{R}\sum_{r=1}^{R} P(\mathbf{\Theta}_i|\mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_{i-1}, \mathbf{\Theta}_{i+1}^{(r)}, \ldots, \mathbf{\Theta}_K^{(r)}, \mathbf{U})$, where $\mathbf{\Theta}_{i+1}^{(r)}, \ldots, \mathbf{\Theta}_K^{(r)}$ are Gibbs samples from $P(\mathbf{\Theta}_{i+1}, \ldots, \mathbf{\Theta}_K|\mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_{i-1}, \mathbf{\Omega})$. Thus, given a range $[K_{min}, K_{max}]$ for the number of common context-aware preferences, we can select the $K \in [K_{min}, K_{max}]$, which maximizes the likelihood.

5. EXPERIMENTAL RESULTS

In this section, we evaluate the performances of the three implementations of the proposed framework based on two kinds of contextual assumptions for predicting user preferences of content categories, namely, CIAP-LDA (**C**ontext conditional **I**ndependency **A**ssumption-based **P**rediction with **LDA** model), CIAP-LDAC (**C**ontext conditional **I**ndependency **A**ssumption-based **P**rediction with **LDAC** model), and CDAP (**C**ontext **D**ependency **A**ssumption-based **P**rediction), with several baseline methods on a real-world dataset.

5.1. Experimental Data

The dataset used in the experiments is collected from many volunteers by a major manufacturer of smart mobile devices. The dataset consists of 8,852,187 context records that contain rich contextual information and usage records of 443 smartphone users spanning from several weeks to several months. Table III shows the concrete types of context data the dataset contains. In the experiments, we classified the 665 unique contents appearing in raw usage records into 12 content categories based on the taxonomy of the Nokia Ovi store (www.ovi.com), which are *Call*, *Web*, *Multimedia*, *Management*, *Game*, *System*, *Navigation*, *Business*, *Reference*, *Social Network Service (SNS)*, *Utility,* and *Others*. In our experiments, we do not utilize the categories *Call* and *Others* because they are not useful for generating corresponding recommendations. Instead, we only utilize the other 10 content categories, which contain 618 unique contents appearing in a total of 408,299 usage records.

Figures 6(a) and 6(b) show the distribution of 618 unique contents in raw usage records with respect to the corresponding content categories and the distribution of context records with respect to the content categories that their corresponding usage

Table III. The Types of Contextual Information in Our Dataset

Context	Value Range
Week	{Monday, Tuesday, . . . , Sunday}
Is a holiday?	{Yes, No}
Day period	{Morning(7:00-11:00), Noon(11:00-14:00), Afternoon(14:00-18:00), Evening(18:00-21:00), Night(21:00-Next day 7:00)}
Time range	{0:00-1:00, 1:00-2:00, . . . ,23:00-24:00}
Profile type	{General, Silent, Meeting, Outdoor, Pager, Offline}
Battery level	{Level 1, Level 2, . . . ,Level 7}
Charging state	{Charging, Complete, Not Connected}
Social location	{Home, Work Place, On the way}

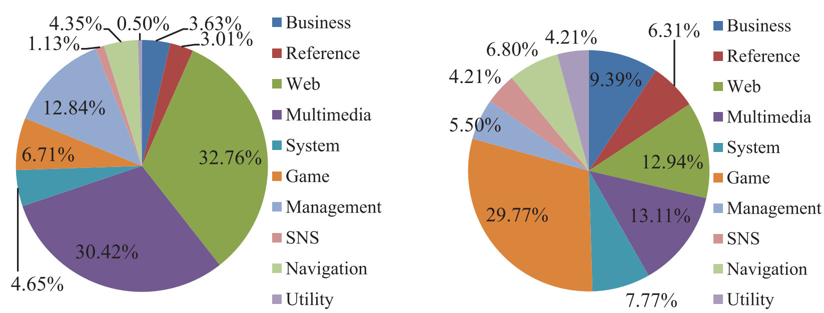

(a) (b)

Fig. 6. (a) The distribution of 618 unique contents w.r.t. the corresponding content categories and (b) the distribution of context records w.r.t. the content categories that their corresponding usage records belong to.

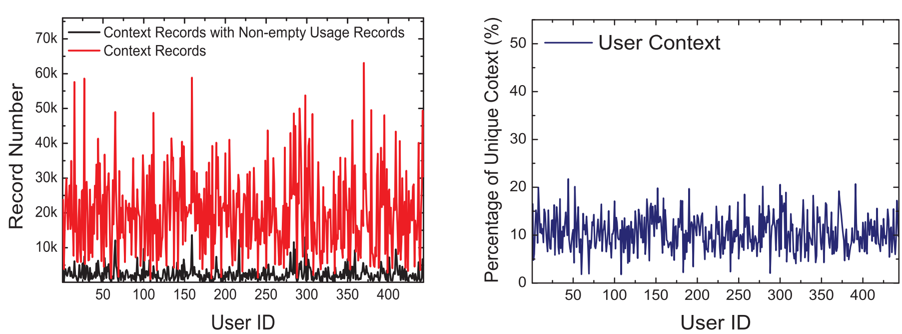

(a) (b)

Fig. 7. (a) The distributions of all context records and the context records with nonempty usage records for all users. (b) The coverage of unique contexts in each user's context log compared with all unique contexts.

records belong to, respectively. The context records with empty usage records are not taken into account.

Figure 7(a) compares the distributions of all context records and the context records with nonempty usage records for all users. From the figure, we can see that usually, although many context records of individual mobile users are collected, only a small proportion of them have nonempty usage records and can be used as training data, which implies the limit of mining personal context-aware preferences only from individual users' context logs. Moreover, Figure 7(b) shows the coverage ratio of unique contexts in each user's context log compared with all unique contexts, from which we can see that the unique contexts in each individual user's context log is relatively limited, which motivates learning from many users' context logs as well.

5.2. Benchmark Methods

First, we select two baseline methods that are transformed from our approaches but only consider individual users' context logs for evaluating whether only considering individual users' context logs is not enough for the personalized context-aware preference mining.

CIAP-i stands for a variant of CIAP that only leverages individual users' context logs. To be specific, in this approach, given a user u and a context C, we predict the user-preferred content categories by ranking each content category c according to $P(c|u, C)$, which can be estimated by $P(c|u, C) \propto \prod_{p \in C} P(c, p|u)$, and $P(c, p|u)$ can be calculated by $P(c, p|u) = \frac{n_{c,p}}{n_{(.)}}$, where $n_{c,p}$ and $n_{(.)}$ indicate the numbers of ACP-feature (c, p) and all ACP-features appearing in the context log of u, respectively.

CDAP-i stands for a variant of CDAP that only leverages individual users' context logs, in which given a user u and a context C, the user-preferred content categories are predicted by ranking each content category c according to $P(c|u, C)$. If C contains some personal content-relevant contexts denoted C^r, where C^r should appear in some behavior patterns of u, we let $P(c|u, C) \propto P(c, C|u) \propto \frac{1}{|C^r_{max}|} \sum_{C^r_{max}} P(c|C^r_{max}, u)$, where C^r_{max} denotes a maximal personal content-relevant subcontext contained by C, and $P(c|C^r_{max}, u)$ can be estimated by the confidence of the association "$C^r_{max} \longrightarrow a$" in the context log of user u. Otherwise, if C does not contain any personal content-relevant subcontext, we do not recommend any content to user u.

Then, we select a state-of-the-art personalized context-aware recommendation approach based on individual users' context logs as a baseline.

CASVM stands for personalized **C**ontext-**A**ware preference prediction by Ranking **SVM**, which is introduced in Kahng et al. [2011]. To be specific, in this approach, given a user's u and a context C, we calculate five types of features $P(c)$, $P(c|u)$, $P(c|p)$, $P(c|u, p)$, and $P(c|u, p_1, \ldots, p_n)$, where p denotes a contextual feature–value pair in C, according to the context log of u, and then leverage Ranking SVM for ranking content category c.

Moreover, to validate the performance of leveraging many users' context logs for mining personal context-aware preferences, we also select two state-of-the-art collaborative filtering (CF)-based approaches as baselines.

CACF stands for **C**ontext-**A**ware preference mining by disjunction **CF**, which is a memory-based CF approach introduced in Lee et al. [2010]. In this approach, the preference score $s_{c,u,C}$ of the user u on content category c in the context C is calculated by the disjunctive aggregation of the estimated preferences of $s_{c,u,p}$, where p is a feature–value pair in C and $s_{c,u,p}$ is measured by counting how many context records of u contain p and contents in category c.

CATF stands for **C**ontext-**A**ware preference mining by **T**ensor **F**actorization, which is a model-based CF approach introduced in Karatzoglou et al. [2010]. In this approach, all users' context logs can be represented by a high-dimensional tensor T, and the objective is to complete the missing values for ranking content categories. Specifically, each value $t_{u,c,v_1,\ldots,v_n}$ in the original tensor is measured by counting how many context records of user u contain contents in category c and feature–value pairs $\{p = (f_i, v_i)\}$.

At last, we select a naive approach that is context aware but not personalized as the last baseline.

CPP stands for **C**ontext-aware **P**opularity-based preference **P**rediction, 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 the user-preferred content categories by the most frequent content categories appearing under C according to all users' historical context logs and then recommend corresponding content categories.

5.3. Evaluation Metrics

In the experiments, we utilize a fivefold cross-validation to evaluate each test approach. To be specific, we first randomly divide each user's context log into five equal parts, and then we use each part as the test data while using the other four parts as the training data in five test rounds. Finally, we report the average performance of the five runs. In the test process, we only take into account the context records with nonempty usage records and use the contexts and the category of the content indicated by the usage record as context inputs and ground truth, respectively. Moreover, to evaluate the ranking of content categories generated by each approach, we leverage two 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 prediction results on the test cases of user u, and $|U|$ indicates the number of test 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 cutoff rank, $P_i(r)$ denotes the precision on the ith test case of u at a given cutoff 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 prediction results. To be specific, $MAR@K = \frac{\sum AR^{(u)}@K}{|U|}$, where $AR^{(u)}@K$ denotes the average recall at top K prediction results on the test cases of user u, and $|U|$ indicates the number of test users. $AR^{(u)}@K$ can be computed by $\frac{1}{N_u}\sum_i \sum_{r=1}^{K} rel_i(r)$.

5.4. Overall Results

According to the parameter estimation approaches introduced in Section 3 and Section 4, the number of common context-aware preferences for LDA, LDAC, and NMF training denoted as K is empirically set to be 15. The further robustness evaluation with varied K will be shown in Section 5.5. For the LDA training, the two parameters α and β are empirically set to be $50/K$ and 0.2 according to Heinrich [2009]. The parameters for the training LDAC model are similar to Bao et al. [2010]. For the Bayesian Matrix Factorization training, according to Schmidt [2009], we use an isotropic noise model and choose a decoupled prior for $\mathbf{\Phi}$ and $\mathbf{\Theta}$ with zero mean $\mu = 0$, and a unit diagonal covariance matrix $\mathbf{\Sigma} = \mathbf{I}$. The maximum iterations of Gibbs sampling are set to be 2,000 in our experiments. Moreover, the behavior patterns are mined by BG-Growth algorithms introduced in Li et al. [2012]. Both our approaches and the baselines are implemented by C++ and the experiments are conducted on a 3GHZ$\times$4 quad-core CPU, 3G main memory PC.

Figure 8 shows the convergence curves of Gibbs sampling for our three implementations of the proposed approach by measuring their log likelihood for the training dataset in one of the five test rounds. From the figure, we can see that the Gibbs sampling of all implementations converges quickly. The convergence curves for other test rounds follow a similar trend. Moreover, each iteration of Gibbs sampling costs on average 89 milliseconds for CIAP-LDA, 125 milliseconds for CIAP-LDAC, and 423 milliseconds for CDAP, respectively. This is because the NMF training is more complex than topic models and the number of associations between context and content category for matrix factorization is greater than the ACP-features in both the LDA and LDAC model.

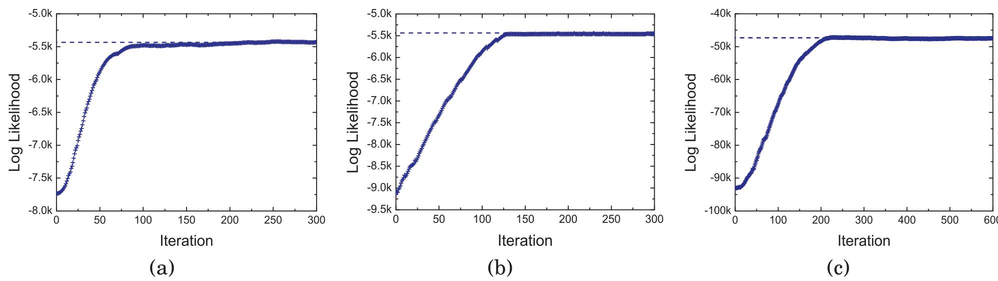

Fig. 8. The log likelihood of training sets in each iteration of Gibbs sampling for (a) CIAP-LDA, (b) CIAP-LDAC, and (c) CDAP.

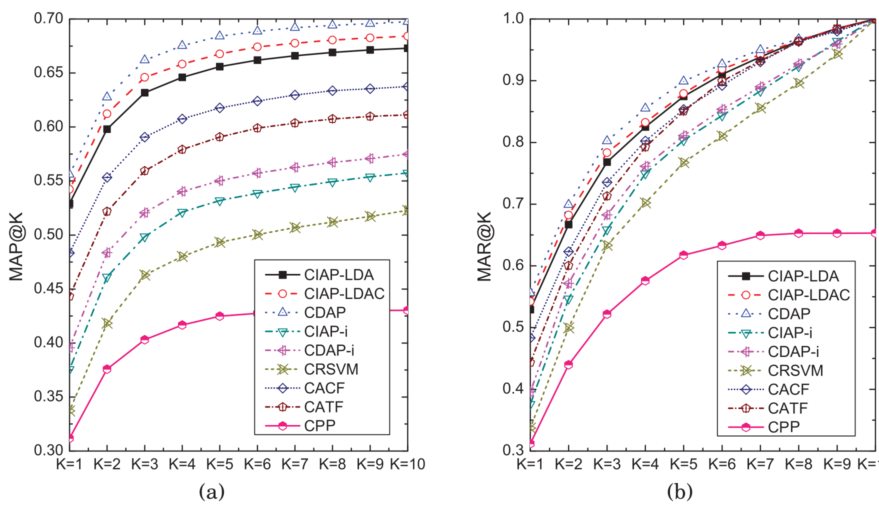

Fig. 9. The average (a) *MAP@K* and (b) *MAR@K* performance of each prediction approach in the fivefold cross-validation.

It is worth noting that neither our approach nor baselines except CATF may be able to generate recommendations for all contexts. To be specific, CIAP-LDA, CIAP-LDAC, CIAP-i, CRSVM, and CACF can only generate preference predictions for the contexts that contain the ACP-features appearing in training sets, CDAP and CDAP-i can only generate preference predictions for the contexts that contain content-relevant subcontexts, and CPP can only predict for the contexts under which there exist some activity records in training sets. In our experiments, we observe that CIAP-LDA, CIAP-LDAC, CIAP-i, CDAP, and CRSVM can cover 100% test contexts, while CDAP-i can cover on average 89.45% test contexts and CPP can cover on average 96.23% test contexts in all five round tests. In the evaluation of these approaches, we only count their *MAP@K* and *MAR@K* for the test cases whose contexts they cover.

We first test the *MAP@K* performance of each test approach with respect to varying K, and the average results in the fivefold cross-validation are shown in Figure 9(a). From the experimental results, we can get some insightful observations as follows. First, our approaches CDAP, CIAP-LDA, and CIAP-LDAC consistently outperform other baselines with varying K, which may indicate that our preference mining framework is more effective for mining context logs than other approaches, such as memory-based and model-based CF methods. Second, CIAP-LDAC outperforms CIAP-LDA slightly, which indicates the effectiveness of the contextual feature priori in LDAC for

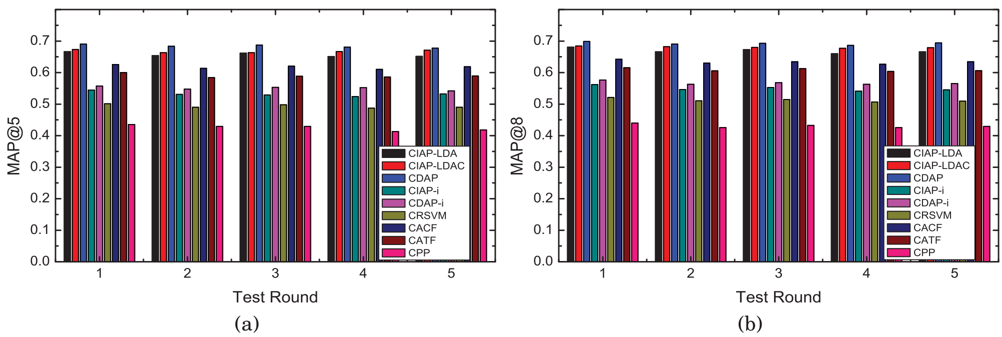

Fig. 10. The $MAP@K$ of each prediction approach when (a) $K = 5$, (b) $K = 8$ in each test round.

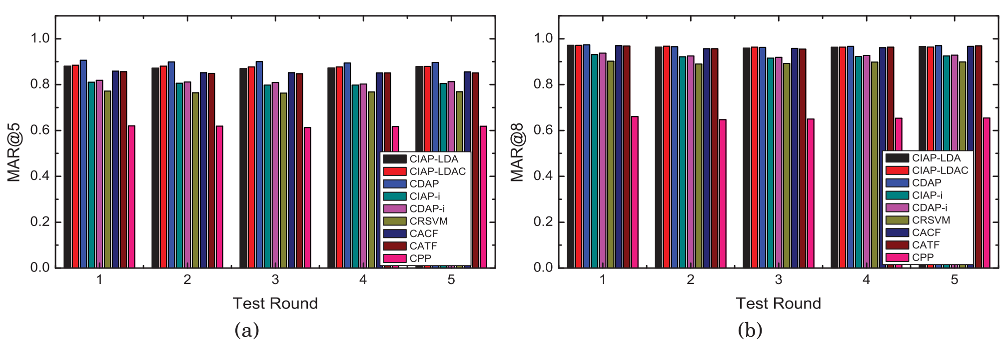

Fig. 11. The $MAR@K$ of each prediction approach when (a) $K = 5$ and (b) $K = 8$ in each test round.

modeling ACP-feature bags. Third, we can see that CDAP outperforms CIAP-LDAC slightly with varying K, which validates that the context dependency assumption is more robust than the context conditional independency assumption. Fourth, the popularity-based approach CPP always has the worst performance in predicting user preferences, which indicates that the popularity-based approach is not suitable for context-aware recommendation. Finally, we also find that two CF-based approaches, CACF and CATF, outperform other individual-based approaches (i.e., CRSVM, CIAP-i, and CDAP-i), which indicates that leveraging many users' context logs other than individual users' context logs can improve the recommendation performance. Figures 10(a) and 10(b) further show the $MAP@K$ of each recommendation approach in each test round when $K = 5$ and $K = 8$, respectively. From the results, we can observe that our approaches consistently outperform other baselines in all five test rounds.

Figure 9(b) shows the average $MAR@K$ of each approach in fivefold cross-validation. From the results, we can observe that it has similar trends to the performance $MAP@K$. Specifically, we find that CIAP-LDA, CIAP-LDAC, and CDAP outperform other baselines and CDAP outperforms CIAP-LDAC slightly in terms of $MAR@K$. Figures 11(a) and 11(b) further show the $MAR@5$ and $MAR@8$ of each recommendation approach in each test round. The results also validate the effectiveness of our approaches in terms of $MAR@K$ performance.

Especially, we conduct a series of paired t-tests with 0.95 confidence level in each K. The test results show that the improvements of CIAP-LDA, CIAP-LDAC, and CDAP on MAP and MAR compared with other baselines are statistically significant. We also study the variances of overall $MAP@K$ and $MAR@K$ of all tested approaches in the fivefold cross-validation for validating the effectiveness of experimental results.

Table IV. The Results of the Mean Deviations of *MAP@K* and *MAR@K* of Each Tested Approach in the Fivefold Cross-Validation

	MAP@1*	MAP@5	MAP@8	MAR@5	MAR@8
CIAP-LDA	9.73×10^{-3}	5.34×10^{-3}	2.96×10^{-3}	1.32×10^{-2}	6.35×10^{-3}
CIAP-LDAC	8.32×10^{-3}	4.13×10^{-3}	2.90×10^{-3}	1.45×10^{-2}	5.97×10^{-3}
CDAP	$\mathbf{7.28 \times 10^{-3}}$	$\mathbf{3.87 \times 10^{-3}}$	$\mathbf{1.04 \times 10^{-3}}$	$\mathbf{8.23 \times 10^{-3}}$	$\mathbf{4.21 \times 10^{-3}}$
CIAP-i	1.52×10^{-2}	9.22×10^{-3}	4.15×10^{-3}	4.43×10^{-2}	8.72×10^{-3}
CDAP-i	1.31×10^{-2}	8.59×10^{-3}	3.87×10^{-3}	4.02×10^{-2}	9.11×10^{-3}
CASVM	2.84×10^{-2}	1.34×10^{-2}	5.44×10^{-3}	5.52×10^{-2}	2.24×10^{-2}
CACF	9.97×10^{-3}	7.01×10^{-3}	3.38×10^{-3}	2.38×10^{-2}	7.98×10^{-3}
CATF	1.27×10^{-2}	8.22×10^{-3}	3.61×10^{-3}	1.74×10^{-2}	7.35×10^{-3}
CPP	4.82×10^{-2}	2.94×10^{-2}	5.89×10^{-3}	7.38×10^{-2}	4.02×10^{-3}

*MAP@1 and MAR@1 have same value for each tested approach.

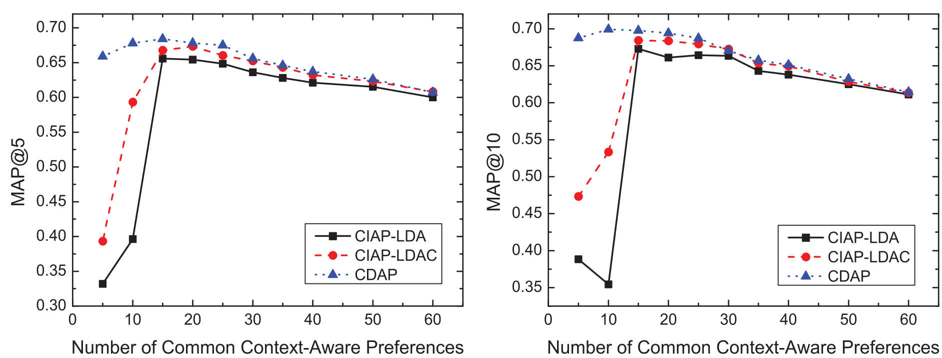

Fig. 12. The (a) *MAP*@5 and (b) *MAP*@10 of CIAP and CDAP with respect to varying number of common context-aware preferences.

Table IV shows the mean deviations of these values of each tested approach in the fivefold cross-validation with $K = 1$, $K = 5$, and $K = 8$. From this table, we can observe that the variances of our approaches (i.e., CIAP-LDA, CIAP-LDAC, CDAP) and the two CF-based approaches (i.e., CACF, CATF) are consistently smaller than other baselines. It implies that individual users' context logs are often very limited, which may influence the stability of recommendation performance.

From these experimental results, we can clearly see that CDAP, CIAP-LDA, and CIAP-LDAC outperform other baselines under different metrics and experimental settings, which demonstrates the effectiveness of our framework for personalized context-aware recommendation. Moreover, CDAP outperforms CIAP-LDA and CIAP-LDAC slightly though its training cost is much higher.

5.5. Robustness Analysis

Since CIAP-LDA, CIAP-LDAC, and CDAP need a parameter to determine the number of common context-aware preferences (i.e., K), Figures 12(a) and 12(b) show the *MAP*@5 and *MAP*@10 of CIAP-LDA, CIAP-LDAC, and CDAP with respect to varying settings of the number. From these figures, we can observe that both *MAP*@5 and *MAP*@10 of CDAP are relatively not sensitive to the parameter. In contrast, the robustness of CIAP-LDA and CIAP-LDAC is not good with small numbers of common context-aware preferences but becomes stable when the setting of the number increases. It may be because CDAP leverages associations between contexts and user content categories for extracting common context-aware preferences and such associations have been

Table V. Prediction Case 1 for Users #162 and #423

Context	{(*Time range: PM21:00-22:00*), (*Profile: Silent*), (*Is holiday: Yes*),(*Day name: Sunday*), (*Day period: Night*), (*Location: Home*)}		
	Top 3 predicted preferences for user #*162*		
Ground truth	***Game***		
CIAP-LDA	**Game** (√)	Multimedia	Web
CIAP-LDAC	**Game** (√)	Multimedia	SNS
CDAP	**Game** (√)	Web	Multimedia
CIAP-i	Multimedia	Business	**Game** (√)
CDAP-i	Web	**Game** (√)	System
CASVM	Multimedia	Web	**Game** (√)
CACF	Multimedia	**Game** (√)	System
CATF	Multimedia	**Game** (√)	Web
CPP	Multimedia	SNS	Web
	Top 3 predicted preferences for user #*423*		
Ground truth	***Web***		
CIAP-LDA	**Web** (√)	Multimedia	SNS
CIAP-LDAC	**Web** (√)	Multimedia	Game
CDAP	**Web** (√)	Reference	Game
CIAP-i	Multimedia	Game	**Web** (√)
CDAP-i	SNS	Multimedia	**Web** (√)
CASVM	Multimedia	Reference	**Web**(√)
CACF	Game	**Web** (√)	Multimedia
CATF	Game	**Web** (√)	System
CPP	Multimedia	SNS	**Web** (√)

filtered from noisy data. Thus, the quality of mined common context-aware preferences is always relatively good with different parameters since the mining is on the basis of pruned training data. In contrast, CIAP-LDA and CIAP-LDAC leverage ACP-features for extracting common context-aware preferences, where ACP-features usually contain more noisy information and thus make the mining results more sensitive to parameters.

5.6. Case Study

In addition to the studies on the overall performance of our approach, we also study the cases in which CIAP-LDA, CIAP-LDAC, and CDAP outperform the baselines.

For example, Table V shows an example of the top three predicted preferences of each approach for two different users given the same context {(*Time range: 21:00-22:00*), (*Profile: Silent*), (*Is holiday: Yes*), (*Day name: Sunday*), (*Day period: Night*), (*Location: Home*)}, which may imply the leisure time at home. Under that context, the ground-truth-preferred content categories of user #162 and user #423 are *Game* and *Web*, respectively. From the results, we can observe that CIAP-LDA, CIAP-LDAC, and CDAP can predict correct content categories for user #162 and user #423 in the top one position. In contrast, CIAP-i, CDAP-i, CASVM, CACF, and CATF predicted correct content categories in lower positions for both users. Moreover, the popularity-based approach CPP always predicts the same preferences for both users and thus does not perform well.

Furthermore, Table VI shows another example of the top three predicted preferences of each approach for user #162 and user #423. The given context is {(*Time range: PM18:00-19:00*), (*Profile: General*), (*Is holiday: No*), (*Day name: Monday*), (*Day period: Evening*), (*Location: on the way*)}, which is similar to the context introduced in our motivating example (i.e., Example 1.1). In this case, the ground truth content categories

Table VI. Prediction Case 2 for Users #162 and #423

Context	{(*Time range: PM18:00-19:00*), (*Profile: General*), (*Is holiday: No*), (*Day name: Monday*), (*Day period: Evening*), (*Location: On the way*)}		
	Top 3 predicted preferences for user #*162*		
Ground truth	***Web***		
CIAP-LDA	**Web** (√)	Multimedia	Game
CIAP-LDAC	**Web** (√)	Multimedia	Game
CDAP	**Web** (√)	SNS	Multimedia
CIAP-i	Game	Navigation	**Web** (√)
CDAP-i	Game	SNS	**Web** (√)
CASVM	Game	Multimedia	**Web**(√)
CACF	Multimedia	**Web** (√)	Game
CATF	Multimedia	**Web** (√)	SNS
CPP	Game	SNS	Multimedia
	Top 3 predicted preferences for user #*423*		
Ground truth	***Multimedia***		
CIAP-LDA	**Multimedia** (√)	Game	Web
CIAP-LDAC	**Multimedia** (√)	Game	Web
CDAP	**Multimedia** (√)	Game	SNS
CIAP-i	Game	Web	**Multimedia** (√)
CDAP-i	Web	SNS	**Multimedia** (√)
CASVM	Web	Reference	**Multimedia**(√)
CACF	Game	**Multimedia** (√)	SNS
CATF	Game	Web	**Multimedia** (√)
CPP	Web	SNS	Game

of user #162 and user #423 are *Web* and *Multimedia*, respectively. From the results, we can observe the similar trend of predicting performance as the last case.

Indeed, our approaches outperform other baselines, especially the two CF-based approaches (i.e, CACF, CATF), because our mining framework can represent personal context-aware preferences of mobile users in a more reasonable way, which can model contexts with respect to different data dependency assumptions. Thus, the three approaches developed based on our framework can be more effective for preference predicting by mining context logs. Moreover, the individual-based approaches (i.e., CRSVM, CIAP-i, and CDAP-i) perform worse in preference predicting than our approaches because individual users' context logs are often very limited, which may influence the stability of recommendation performance.

Finally, to further study the reason that our approaches can outperform other baselines, we investigate the mined personal context-aware preferences for user #162. To be specific, in our framework, the personal preference of each user can be represented as a distribution of these common context-aware preferences. In fact, the common context-aware preference z^* that a user u has the highest probability of belonging to, that is, $z^* = \arg\max_z P(z|u)$, will have the most influential impact on the personal preference of u. Therefore, for user #162, we manually inspect the most important common preference z^* and study whether it contains some relationships between the given context and ground truth content. Tables VII to IX show the different z^* mined by CIAP-LDA, CIAP-LDAC, and CDAP for user #162, respectively. Note that the common preferences mined by both CIAP-LDA and CIAP-LDAC are constituted by ACP-features, while the common preferences mined by CDAP are constituted by behavior patterns (limited by space, we only show the top 10 ACP-features and behavior patterns in each common preference). From these tables, we can observe that there are many ACP-features and behavior patterns (i.e., labeled in bold) that have explicit relationships with the given

Table VII. The Common Preference z^* Mined by CIAP-LDA for User #162

1	**(Is holiday?: Yes) → Game**
2	(Is holiday?: Yes) → Multimedia
3	(Day period: Evening) → Game
4	(Location: Home)→ Web
5	**(Location: Home) → Game**
6	**(Profile: Silent) → Game**
7	**(Profile: General) → Web**
8	**(Day period: Night) → Game**
9	(Location: Home) → Multimedia
10	**(Day period: Evening) → Web**

Table VIII. The Common Preference z^* Mined by CIAP-LDAC for User #162

1	**(Is holiday?: Yes) → Game**
2	**(Location: Home) → Game**
3	**(Is holiday?: No) → Web**
4	(Location: Home) → SNS
5	(Location: Home)→ Web
6	**(Profile: Silent) → Game**
7	**(Day period: Night) → Game**
8	**(Profile: General) → Web**
9	(Day name: Sunday) → Web
10	**(Location: On the way) → Web**

Table IX. The Common Preference z^* Mined by CDAP for User #162

1	**(Is holiday?: Yes) (Location: Home) → Game**
2	(Location: Home) (Profile: General) → Web
3	(Is holiday?: Yes) (Location: Home)→ Multimedia
4	**(Profile: Silent) (Time period: Night) → Game**
5	**(Location: Home) (Profile: Silent) → Game**
6	**(Location: Home) (Day Period: Night) → Game**
7	**(Is holiday?: No) (Day Period: Evening) → Web**
8	(Profile: General) (Location: home)→ Game
9	(Day name: Sunday) (Profile: General) → Web
10	**(Location: On the way) (Day Period: Evening) → Web**

context and ground truth content category. Therefore, our approaches can predict the corresponding content category with high ranking in the results. In conclusion, these case studies clearly validate the effectiveness of our preference mining framework and approaches.

6. RELATED WORK

Many previous works about personalized context-aware recommendation for mobile users have been reported. For example, Tung and Soo [2004] have proposed a prototype design for building a personalized recommender system to recommend travel-related information according to users' contextual information. Park et al. [2007] proposed a location-based personalized recommender system that can reflect users' personal preferences by modeling user contextual information through Bayesian networks. Bader et al. [2011] proposed a novel context-aware approach to recommending points-of-interest (POIs) for users in an automotive scenario. Specifically, they studied the scenario of recommending gas stations for car drivers by leveraging Multi-Criteria Decision Making (MCDM)-based methods to modeling context and different routes. However, most of these works only leverage individual users' 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. [2007] proposed a hybrid framework named "play.tools" for recommending mobile applications by leveraging users' context information. This recommendation framework is based on what other users have installed in similar contexts and will be liked by a given user. Kim et al. [2010] 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. [2010] 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 three-dimensional matrix, namely, tensor, and perform tensor factorization with several constraints to capture users' preferences. Karatzoglou et al. [2010] 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, we usually cannot obtain such rating data in user mobile devices. In contrast, it is relatively easier to collect context logs, which contain the users' historical context information and their usage records, which can be used for mining context-aware user preferences.

Recently, some researchers studied how to mine event logs for personalized context-aware recommendation. For example, Lee et al. [2010] investigated converting the event logs into implicit ratings and tested various memory-based CF approaches to make personalized context-aware recommendation. Kahng et al. [2011] proposed a novel approach for ranking items in event logs for personalized context-aware recommendation. Compared with context logs, event logs only record the context records with nonempty usage records and thus lose some discriminative information to capture the relevance between contexts and content categories [Cao et al. 2010]. But these event-log-based approaches can still be used for mining context logs toward personalized context-aware recommendation. Recently, Yu et al. [2012] proposed a novel personalized context-aware mobile recommender system by analyzing mobile users' context logs. The proposed approach is based on the Latent Dirichlet Allocation topic model and scalable for multiple contextual features. However, this approach only focuses on the situation of context-independent assumption, which is similar to our CIAP-LDA approach. Different from these works, we propose a novel framework that can explicitly model common context-aware preferences and represent individual users' personal context-aware preferences by distributions of common context-aware preferences with different context dependency assumptions. Moreover, one proposed approach in our framework (CDAP) can utilize the context records with only empty usage records and context-dependent assumption.

7. CONCLUDING REMARKS

In this article, we proposed to exploit user context logs for mining the personal context-aware preferences of mobile users. First, we identified common context-aware preferences from the context logs of many users. Then, the personal context-aware preference of an individual user can be represented as a distribution of common context-aware preferences. Moreover, we designed two methods for mining common context-aware preferences based on two different assumptions about context data dependency. Finally, the experimental results from a real-world dataset clearly showed that the proposed approach could achieve better performances than benchmark methods for mining personal context-aware preferences, and the one implementation based on the independent assumption of context data slightly outperforms another one but has higher computational cost.

As mentioned earlier, the proposed user context-aware preference mining approaches are based on the analysis of many users' daily context logs, which will be very huge in quantity. Therefore, how to efficiently store and exploit our mining approaches on such "big data" are very challenging problems to be further studied. Moreover, how to capture users' preference drifts and protect their privacy are also valuable research aspects in our future plan. Last but not least, in the future, we will try to integrate our approaches with various real-world services, such as mobile context-aware advertising, for enhancing user experience.

A. APPENDIX: GIBBS SAMPLING DERIVATION FOR CIAP-LDAC APPROACH

In this article, we leverage the Gibbs sampling-based approach introduced in Bao et al. [2010] to train the LDAC model in our CIAP-LDAC approach. To be specific, this method begins with a random assignment of common context-aware preferences 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 common context-aware preference to each ACP-feature, which is conditional on the assignment of all other ACP-features. Then a new assignment of common context-aware preferences 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 common context-aware preference.

According to the introduction in Section 3.1.2 and the graphical representation in Figure 4(b), we have

$$\begin{aligned} P(d, \theta, z_d, \pi_d, \Phi|\alpha, \beta, \gamma) &= P(\theta_d|\alpha)P(\Phi|\beta)P(\pi_d|\gamma) \\ &\quad \times \left(\prod_{i=1}^{N_d} P(c_i, v_{d,i}|z_{d,i}, f_{d,i}, \Phi)P(f_{d,i}|\pi_d)P(z_{d,i}|\theta_d)\right), \end{aligned}$$

where $\Phi = \{\phi_{k,f}\}$, $z_d = \{z_{d,i}\}$, and N_d denotes the number of ACP-features in the ACP-feature bag d. If we denote the token (d, i) as m, the Gibbs sampler of common context-aware preference z_m in each sampling round can be computed as follows:

$$\begin{aligned} P(z_m = k|Z_{\neg m}, D) &\propto P(z_m = k, Z_{\neg m}, D) \\ &\propto P(c_i, v_m|z_m = k, Z_{\neg m}, F, V_{\neg m}) \\ &\quad \times P(z_m = k|Z_{\neg m}), \end{aligned}$$

where $\neg m$ means removing the contextual feature–value pair $(f_m : v_m)$ from corpus D, and Z is the labels of the common context-aware preferences in D. Furthermore, we have the following estimation:

$$\begin{aligned} P(c_i, v_m|z_m = k, Z_{\neg m}, F, V_{\neg m}) &= \frac{n_{\neg m,k,f_m,c_i,v_m} + \beta_{c_i,v_m}}{\sum_v n_{\neg m,k,f_m,c_i,v} + \sum_{v \in V_{f_m}} \beta_v} \\ P(z_m = k|Z_{\neg m}) &= \frac{n_{d,\neg m,k} + \alpha_k}{\sum_{k'=1}^{K} n_{d,\neg m,k'} + \sum_{k'=1}^{K} \alpha_{k'}}, \end{aligned}$$

where $n_{\neg m,k,f,c,v}$ indicates the frequency that the ACP-feature (c, p) $(p = (f : v))$ is labeled with the kth common context-aware preference in all ACP-feature bags after removing the mth ACP-feature, and $n_{d,\neg m,k}$ indicates the number of ACP-features labeled with the kth common context-aware preferences in d except for the mth one.

After the training process, we can obtain the probabilities $P(z|u)$ by Equation (1) and $P(c, p|z)$ as follows:

$$P(c, p|z_k) = P(c, v_p|f_p, z)P(f_p), \tag{6}$$

where $p = (f_p : v_p)$ and

$$P(c, v_p | f_p, z) = \frac{n_{z,c,f_p,v_p} + \beta_{c,v_p}}{\sum_v n_{z,c,f_p,v} + \sum_{v \in V_{fm}} \beta_v},$$

$$P(f_p) = \frac{\sum_z \sum_v n_{k,c,f_p,v} + \gamma_{f_m}}{\sum_f \sum_z \sum_v n_{z,c,f_p,v} + \sum_f \gamma_f},$$

where n_{z,c,f_p,v_p} is the frequency of ACP-feature (c, p) that has been assigned to common context-aware preference z, and V_{f_m} is the number of values for contextual feature f_m.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation for Distinguished Young Scholars of China under Grant 61325010, in part by the Natural Science Foundation of China under Grant 71329201, in part by the National High Technology Research and Development Program of China under Grant 2014AA015203, in part by the National Science Foundation under Grant CCF-1018151 and Grant IIS-1256016.

REFERENCES

Rakesh Agrawal and Ramakrishnan Srikant. 1994. Fast algorithms for mining association rules in large databases. In *Proceedings of the 20th International Conference on Very Large Data Bases (VLDB'94)*. Morgan Kaufmann, San Francisco, CA, 487–499.

Leif Azzopardi, Mark Girolami, and Keith van Risjbergen. 2003. Investigating the relationship between language model perplexity and IR precision-recall measures. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'03)*. ACM, New York, NY, 369–370. DOI:http://dx.doi.org/10.1145/860435.860505

Roland Bader, Eugen Neufeld, Wolfgang Woerndl, and Vivian Prinz. 2011. Context-aware POI recommendations in an automotive scenario using multi-criteria decision making methods. In *Proceedings of the 2011 Workshop on Context-Awareness in Retrieval and Recommendation (CaRR'11)*. ACM, New York, NY, 23–30. DOI:http://dx.doi.org/10.1145/1961634.1961640

Tengfei Bao, Happia Cao, Enhong Chen, Jilei Tian, and Hui Xiong. 2010. An unsupervised approach to modeling personalized contexts of mobile users. In *Proceedings of the 2010 IEEE International Conference on Data Mining (ICDM'10)*. IEEE Computer Society, Washington, DC, 38–47. DOI:http://dx.doi.org/10.1109/ICDM.2010.16

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Lantent Dirichlet allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.

Huanhuan Cao, Tengfei Bao, Qiang Yang, Enhong Chen, and Jilei Tian. 2010. An effective approach for mining mobile user habits. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM'10)*. ACM, New York, NY, 1677–1680. DOI:http://dx.doi.org/10.1145/1871437.1871702

Nathan Eagle, Aaron Clauset, and John A. Quinn. 2009. Location segmentation, inference and prediction for anticipatory computing. In *AAAI Spring Symposium on Technosocial Predictive Analytics*.

Andreas Emrich, Alexandra Chapko, and Dirk Werth. 2009. Context-aware recommendations on mobile services: The m:Ciudad approach. In *Proceedings of the 4th European Conference on Smart Sensing and Context (EuroSSC'09)*. Springer-Verlag, Berlin, 107–120.

Yong Ge, Qi Liu, Hui Xiong, Alexander Tuzhilin, and Jian Chen. 2011. Cost-aware travel tour recommendation. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'11)*. ACM, New York, NY, 983–991. DOI:http://dx.doi.org/10.1145/2020408.2020568

Thomas L. Griffiths and Mark Steyvers. 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America* 101, suppl1 (2004), 5228–5235.

Gregor Heinrich. 2005. *Parameter Estimation for text Analysis*. Technical report, University of Lipzig (2005).

Kyoung Jae Kim, Hyunchul Ahn, and Sangwon Jeong. 2010. Context-aware recommender systems using data mining techniques. *Journal of World Academy of Science, Engineering and Technology* 64 (2010), 357–362.

Minsuk Kahng, Sangkeun Lee, and Sang-goo Lee. 2011. Ranking in context-aware recommender systems. In *Proceedings of the 20th International Conference Companion on World Wide Web (WWW'11)*. ACM, New York, NY, 65–66. DOI:http://dx.doi.org/10.1145/1963192.1963226

Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. 2010. Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys'10)*. ACM, New York, NY, 79–86. DOI:http://dx.doi.org/10.1145/1864708.1864727

Jennifer R. Kwapisz, Gary M. Weiss, and Samuel A. Moore. 2011. Activity recognition using cell phone accelerometers. *SIGKDD Explorer Newsletter* 12, 2 (2011), 74–82.

Dongjoo Lee, Sung Park, Minsuk Kahng, Sangkeun Lee, and Sang-goo Lee. 2010. Exploiting contextual information from event logs for personalized recommendation. *Computer and Information Science* 121–139.

Xueying Li, Huanhuan Cao, Hui Xiong, Enhong Chen, and Jilei Tian. 2012. BP-Growth: Searching strategies for efficient behavior pattern mining. In *Proceedings of 13th IEEE International Conference on Mobile Data Management (MDM'12)*.

Qi Liu, Yong Ge, Zhongmou Li, Enhong Chen, and Hui Xiong. 2011. Personalized travel package recommendation. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining (ICDM'11)*. IEEE Computer Society, Washington, DC, 407–416. DOI:http://dx.doi.org/10.1109/ICDM.2011.118

Qi Liu, Haiping Ma, Enhong Chen, and Hui Xiong. 2013. A survey of context-aware mobile recommendations. *International Journal of Information Technology and Decision Making* 12, 1 (2013), 139–172. DOI:http://dx.doi.org/10.1142/S0219622013500077

Moon-Hee Park, Jin-Hyuk Hong, and Sung-Bae Cho. 2007. Location-based recommendation system using bayesian user s preference model in mobile devices. In *Ubiquitous Intelligence and Computing, 4th International Conference, UIC 2007, Hong Kong, China, July 11-13, 2007, Proceedings (Lecture Notes in Computer Science)*, Jadwiga Indulska, Jianhua Ma, Laurence Tianruo Yang, Theo Ungerer, and Jiannong Cao (Eds.), Vol. 4611. Springer, 1130–1139. DOI:http://dx.doi.org/10.1007/978-3-540-73549-6_110

Ivana Podnar, Manfred Hauswirth, and Mehdi Jazayeri. 2002. Mobile push: Delivering content to mobile users. In *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCSW'02)*. IEEE Computer Society, Washington, DC, 563–570.

Mikkel N. Schmidt. 2009. Linearly constrained Bayesian matrix factorization for blind source separation. In *Proceedings of the 23rd Annual Conference on Neural Information Processing Systems*. 1624–1632.

Mikkel N. Schmidt, Ole Winther, and Lars Kai Hansen. 2009. Bayesian non-negative matrix factorization. In *Proceedings of 2009 International Conference on Independent Component Analysis and Signal Separation*. 540–547.

Timothy Sohn, Kevin A. Li, Gunny Lee, Ian Smith, James Scott, and William G. Griswold. 2005. Place-Its: A study of location-based reminders on mobile phones. In *Proceedings of the 5th International Conference on Ubiquitous Computing (Ubicomp'05)*. 232–250.

Hung-Wen Tung and Von-Wun Soo. 2004. A personalized restaurant recommender agent for mobile e-service. In *Proceedings of the 2004 IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE'04)*. IEEE Computer Society, Washington, DC, 259–262.

Wolfgang Woerndl, Christian Schueller, and Rolf Wojtech. 2007. A hybrid recommender system for context-aware recommendations of mobile applications. In *Proceedings of the 2007 IEEE 23rd International Conference on Data Engineering Workshop (ICDEW'07)*. IEEE Computer Society, Washington, DC, 871–878. DOI:http://dx.doi.org/10.1109/ICDEW.2007.4401078

Kuifei Yu, Baoxian Zhang, Hengshu Zhu, Huanhuan Cao, and Jilei Tian. 2012. Towards personalized context-aware recommendation by mining context logs through topic models. In *Proceedings of the 16th Pacific-Asia conference on Advances in Knowledge Discovery and Data Mining - Volume Part I (PAKDD'12)*. Springer-Verlag, Berlin, 431–443. DOI:http://dx.doi.org/10.1007/978-3-642-30217-6_36

Vincent Wenchen Zheng, Bin Cao, Yu Zheng, Xing Xie, and Qiang Yang. 2010. Collaborative filtering meets mobile recommendation: A user-centered approach. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI'10)*. 236–241.

Yu Zheng, Lizhu Zhang, Xing Xie, and Wei-Ying Ma. 2009. Mining interesting locations and travel sequences from GPS trajectories. In *Proceedings of the 18th International Conference on World Wide Web (WWW'09)*. ACM, New York, NY, 791–800. DOI:http://dx.doi.org/10.1145/1526709.1526816

Hengshu Zhu, Huanhuan Cao, Enhong Chen, Hui Xiong, and Jilei Tian. 2012a. Exploiting enriched contextual information for mobile app classification. In *Proceedings of the 21st ACM International*

Conference on Information and Knowledge Management (CIKM'12). ACM, New York, NY, 1617–1621. DOI:http://dx.doi.org/10.1145/2396761.2398484

Hengshu Zhu, Enhong Chen, Kuifei Yu, Huanhuan Cao, Hui Xiong, and Jilei Tian. 2012b. Mining personal context-aware preferences for mobile users. In *Proceedings of the 2012 IEEE 12th International Conference on Data Mining (ICDM'12)*. IEEE Computer Society, Washington, DC, 1212–1217. DOI:http://dx.doi.org/10.1109/ICDM.2012.31

Received November 2012; revised June 2013; accepted September 2013