On-Device Mining of Mobile Users' Co-Occurrence Patterns

Vijay Srinivasan, Saeed Moghaddam, Abhishek Mukherji, Kiran K. Rachuri, Chenren Xu*, Emmanuel Munguia Tapia
Samsung Research America, San Jose, CA, USA

1. Introduction

Knowing the user's longitudinal behavioral patterns and mobile device usage patterns can enable exciting new context-aware applications and services. For the example user in Figure 1, based on the user's typical daily patterns: (i) we can preload news apps in the morning to improve user experience by reducing loading time, (ii) we can provide task shortcuts on mobile devices based on the user's typical actions when at work in the afternoon, (iii) we can create automatic reminders to charge the phone just before the user typically goes to sleep.

Most existing research on frequent pattern mining has been applied to data processed in the cloud or dedicated high-power desktop environments [1, 2]. In this work, we develop an ***on-device mobile mining*** service. A key advantage of on-device mining is that the user's frequent patterns are mined entirely on the phone without any of the privacy concerns or other costs associated with uploading private, personal data to the cloud. However, a key challenge is that the on-device pattern mining must have a low resource overhead in order to avoid interfering with user-facing applications or services on mobile devices.

2. Co-Occurrence Engine

We present our ongoing work and results in the design and development of an on-device mobile **context co-occurrence engine** for mining frequent patterns over context events that occur *together* around the same time or location. **Context events** include the user's location, physical activities, and usage of the phone's hardware and software. Our goal is to provide easily understandable frequent patterns in the form of **association rules** [1] for use in a number of behavior-aware mobile applications and services. For example the rule: Call Alex ← (AtHome, Time=9-10 PM), indicates that the user typically calls Alex when at home between 9 to 10 PM.

Figure 2 shows the overview of our co-occurrence engine system running on the mobile device. Our **MobileMiner** algorithm accepts as input a rich dataset consisting of timestamped, mobile **context** data. The **Co-Occurrence Itemset Extractor** first extracts itemsets (or baskets as they are typically referred to in *market basket mining* [1]) of context events which occur together in the input mobile context data. An example itemset is: {Time = 9 - 10AM, At_Home, Reading News, Weather=Cloudy}. The extractor then performs a compression step where several duplicate itemsets that occur repeatedly are compressed into a single unique itemset that is weighted by the total duration of co-occurrence of the itemset over the context data. Our **WeMit** (**W**eighted **E**fficient **Mi**ning of **T**emporal Patterns) **Itemset Miner** mines for frequent mobile user patterns on smartphones by leveraging the duration weights of input itemsets to achieve a faster running time compared to the standard Apriori algorithm [2]. The **Context-Aware Scheduler** intelligently schedules the co-occurrence pattern mining so that routine user activities on the phone are not affected while the **Pattern Query Interface** allows external client applications to retrieve patterns for their use.

Table 1 shows five example association rules generated for two sample users from over 27 days of context data along with the **confidence** and the user for which the pattern was created; for the rule A←B, **confidence** is defined as the conditional probability P(A|B). For example, in rule 1, Alex is a friend that user 1 typically calls on his drive home from work. Rule 2 indicates that if user 1 forgets to charge his phone before sleeping, battery level is low in the mornings. From rules 3, 4, and 5 in Figure 3, we see respectively that user 2: (1) does not charge his phone at work, (2) is typically at home at night around 10 PM, and (3) typically calls home around 6 PM.

Figure 1: Example user patterns.

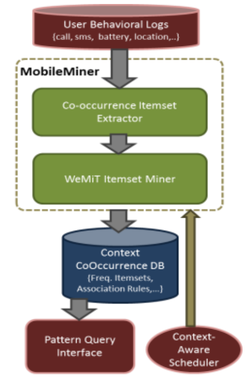

Figure 2: Co-occurrence engine architecture.

Rule No.	Mined user pattern	Conf-idence	User Id
1	Call Alex ← Outgoing call, Time of day=6 to 7pm, Outside, Moving	100 %	1
2	Battery < 20% ← Time of day=6 to 7pm, Phone not charging	100 %	1
3	Phone not charging ← At work	91.9%	2
4	At Home ← Time of day =10 to 11pm, Stationary	100 %	2
5	Call home ← Outgoing call, Time of day=6 to 7pm, Battery = 40 to 60%, Phone not charging	100 %	2

Table 1: Example co-occurrence patterns

3. Ongoing work

In ongoing work, we are evaluating the performance and comparing our co-occurrence engine with standard data mining algorithms over data collected from 106 subjects over 2-5 months. We are also implementing and performing field evaluations of several smartphone services that leverage our co-occurrence patterns, including smart content pre-loading, automatic UI shortcuts for making calls or launching apps, and smart reminders for mobile battery life management.

4. References

[1] Nath, Suman. "ACE: exploiting correlation for energy-efficient and continuous context sensing" In MobiSys 2012 pages 29—42.

[2] Agrawal, Rakesh, and Ramakrishnan Srikant. "Fast algorithms for mining association rules." In VLDB 1994.

* Chenren Xu worked on this paper as an intern