# Understanding Usage States on Mobile Devices

**Chakajkla Jesdabodi**
University of Hamburg
Hamburg, Germany
jesdabodi@informatik.uni-hamburg.de

**Walid Maalej**
University of Hamburg
Hamburg, Germany
maalej@informatik.uni-hamburg.de

## ABSTRACT
Nowadays, mobile apps are used for nearly every situation: for planning the day, communicating with colleagues, ordering goods, or entertaining and socializing. To understand users expectations in each situation and to provide context-aware services, researchers and app vendors started to capture users' interaction with the smartphone and to model user's behavior. This paper reports on a behavioral study based on app usage data logged over one year and the corresponding apps descriptions from the app store. Using Topic Modeling and clustering techniques, we segmented the usage data into meaningful clusters that correspond to different "states", in which users normally use their smartphone, e.g. socializing or consuming media. Researchers and app-vendors can use the insights from our work to improve their contextual recommendation techniques and the overall usage experience.

### ACM Classification Keywords
H.3.4 Systems and Software: User profiles and alert services

### Author Keywords
Apps; Usage Data; Behavioral Profiles; Intent Identification

## INTRODUCTION
Mobile devices are increasingly becoming the main platform that people use to perform everydays virtual and physical activities such as browsing the web, messaging, navigating, or jogging. Numerous apps that allow users to perform these activities can easily be accessed and deployed via distribution platforms such as the Google Play or the Apple App Store. The main goal of this research is to examine whether the current user activity (i.e. usage state) such as planning a workday or consuming media can be derived from the sequence of app usages (i.e. app usage trace) that is often captured on the mobile device. We focus on the following research questions:

1. Can we identify distinctive usage states of mobile device users from their app usage traces?
2. How do usage states and their patterns differ within and between the users?


*UbiComp '15*, September 711, 2015, Osaka, Japan.

ACM 978-1-4503-3574-4/15/09...$15.00.
http://dx.doi.org/10.1145/2750858.2805837

Answering these questions would have impact on building predictive services and on improving the user experience. For example, it is useful to know whether a user is working, gaming, or shopping before recommending certain information or certain actions. Intelligent personal assistant like Siri and Google Now as well as content or product recommendation apps like Facebook or Amazon can use this information to re-rank their recommendation items. Finally, knowing and understanding the usage state and recurrent user behavior can help to understand the user needs, better interpret their feedback, and capture their requirements [11, 12].

We report on a behavioral study to answer the research question based on the LiveLab dataset [15], which was collected from 24 iPhone users between February 2010 and April 2011. The data contains about 1.1 Million accumulated app usages of 2325 unique apps corresponding to 20 App Store categories (all but Catalogs, Newsstand, and Food & Drink).

We applied a *sessionization* method of the usage data based on pause time instead of relying on the screen switch on/off events. We then extended the usage traces with the *official apps descriptions* from the Apple Store and created a Topic Model [2] for the apps – extracting meaningful topics of the usage sessions, which correspond to the apps functionality described in the app store. We identified a list of 27 concrete app topics for mobile usage. We then used a Frequent Itemset mining algorithm to find the number of underlying usage states and clustered the usage data with K-Means algorithm. The clusters from the K-Means algorithm represent the behavioral states, each with a different topics distribution.

## USAGE STATE GENERATION

### Trace Sessionization
We performed a time-based sessionization of the usage traces into usage sessions. Each session $S$ can have multiple apps and last for a certain duration. For each user, we calculated the mean pause time between apps $timeThreshold$ and its standard deviation $timeStdev$. This pause time reflects the time user switches between two apps. It is detected by taking the mean of the home screen pause below 2 minutes. This time-based sessionization detects whether a user is switching to another app or ending a usage session.

We formally define a session as a tuple of apps $a_i$ in $Ap(u)$ the set of installed apps on user's $u$ mobile device:

$$S_k = (a_i | a_i \in Ap(u)), \Delta_{i,j} t < timeThreshold + (2.33 * timeStdev) \quad (1)$$

$k$ is the number of apps in the session. $\Delta_{i,j}t$ represents the time gap between two consecutive apps in $S$. 2.33 is the Z-critical value for one-tailed test at a significance level of 0.01.

### Extracting Usage States

A usage state is a repetitive state in the behavioral profile of a user that corresponds to a certain recurrent situation or activity of the user as a part of her tasks. A user performs a certain task in one or more sessions using specific apps. In order to mine different usage states, we need a way to map a usage session to a meaningful representation of usage state.

With **Latent Dirichlet Allocation** (LDA) [2], we can assign a given app to multiple topics that reveal what users typically do with the app. For example, WhatsApp can be used for both social networking and messaging. App descriptions can be edited by the vendors to report on the functionality and new features. Using topics to represent usage states is more flexible and versatile than using app names. First, one can find thousands of apps that can be used for the same task or reflect the same usage state. Second, app names can change but the usage states, activities, and tasks remain the same.

Furthermore, topics present several advantages over using the official app categories from the app store. First, app topics derived from the descriptions are applicable across different mobile platforms, while the categories slightly differ between the stores. Second, app store category can be incorrectly assigned by developers for optimum ranking, e.g. users are more willing to pay for apps in a lifestyle category than in social networking. Third, one app can have several topics, while most store allow only one category for an app. Finally, users may use several apps for the same tasks.

We used the Apple AppStore API to retrieve the descriptions and category for each app in the dataset. We generated a document-topic matrix using Mallet tool [13]. Each rows of the matrix represents an app, and the probability distribution of the topics for this app. We varied the number of LDA topics to be generated from 25 to 32 and manually checked the coherence of the output. We grouped gaming apps into one topic and generated a total of 27 topics. Table 1 shows the topics, the top keywords, and apps for each topic.

We validated the topics based on 9 randomly selected apps. For each app, we retrieved the set of topics which has a distribution above the LDA threshold of 0.20 and compared the automated topic assignments with a manual labeling for the apps. We then computed the precision and recall. Precision is the fraction of the correctly assigned topics divided by the total number of assignments. Recall is the fraction of the correctly assigned topics divided by the sum of correct assignments and the number of non-assigned topics (false negatives). For these 9 apps, we obtained an average precision of 81.5% and recall of 67%.

We then used **Frequent Itemset Mining** [1] to discover the possible number of states for a user given his usage traces. The items within an Itemset are one of the 27 topics in Table 1 and are derived from the apps used in one session. In the Itemset mining process, we set the LDA probability threshold to 0.2. This allows for an app to belong to several topics.

| # | Top keywords | Topic | Example apps |
|---|---|---|---|
| *1* | *entertain, voic, sport* | *entertaining/sports* | *WorldCup,SportsTap* |
| *2* | *photo, video, camera* | *photography* | *Photos,ShutterFly* |
| *3* | *note, task, list* | *task planning* | *Notes,iStudent* |
| *4* | *util, calcul, unit* | *iOS utilities* | *Calculator, iDisk* |
| *5* | *educ, student, learn* | *learning* | *BrainTeaser, SkyGazer* |
| *6* | *languag, word, dictionari* | *language* | *MultiMagicDic,Theasaurus* |
| *7* | *book, read, comic* | *reading* | *Nook,LovePoems* |
| *8* | *lifestyl, love, dai* | *lifestyle* | *CraigList,DealMap* |
| *9* | *weather, forecast, result* | *weather* | *Weather,TWC* |
| *10* | *social, friend, share* | *socializing* | *Meebo,BrightKite* |
| *11* | *featur, time, iphon* | *support* | *AppStore* |
| *12* | *travel, navig, map* | *navigating* | *Maps,GPSLite* |
| *13* | *iphon, ipod, touch* | *device feature* | *Preference, Timer* |
| *14* | *entertain, music, photographi* | *relaxing* | *Pandora,Youtube* |
| *15* | *fit, health, track* | *fitness* | *Nike,Babymed* |
| *16* | *subscript, call, account* | *phone-related* | *Phone, TextPlus* |
| *17* | *prayer, bibl, vers* | *praying* | *HolyBible, LifeChurch* |
| *18* | *health, medic, fit* | *health* | *CardioMath, 3D4Medical* |
| *19* | *product, file, document* | *working* | *eToDo, iSpreadSheet* |
| *20* | *send, messag, contact* | *communicating* | *SMS,Mail* |
| *21* | *new, video, watch, brows* | *browsing news* | *Newstand, Safari* |
| *22* | *music, record, sound* | *music apps* | *DigiDrummer,GuitarToolKit* |
| *23* | *financ, account, card, credit* | *finance* | *BankOfAmerica,CitiMobile* |
| *24* | *food, drink, recip* | *eating* | *RecipeFinder, CookBook* |
| *25* | *store, shop, product* | *shopping* | *Amazon, Coupons* |
| *26* | *sleep, sound, alarm, clock* | *sleeping* | *SleepMachine,iNightClock* |
| *27* | *game, play, fun* | *gaming* | *WordsWithFriend, DoodleJump* |

**Table 1. App topics with top keywords and example apps. Topics ordering is generated randomly by the LDA process. The keywords are ranked based on their probability within a topic.**

We set the minimum support $minSup$ in the Apriori algorithm to 0.05 to discover more itemsets that may represent a usage state. The parameters are fine-tuned iteratively based on the dataset. The number of frequent itemsets is used to set the number of clusters in the **K-means clustering**. Each usage session generates a feature vector that is used by the clustering. The vector has a fixed dimension of 27, which corresponds to the total number of topics. We define a feature vector $\bar{x}_{S_k}$ of a given usage session $S$ is as follows:

$$\bar{x}_{S_k} = [v_1, ..., v_{27}]\,,\; v_i = \sum_{n=1}^{k}(d_{top\,i,n} * p_{top\,i,n}) * score_{top\,i} \qquad (2)$$

$d_{top\,i,n}$ and $p_{top\,i,n}$ is the usage duration and the probability of app $n$ for belonging to a topic $top\,i$ respectively. $score_{top\ i}$ is the inverse-document-frequency $idf$ of $top\,i$ with respect to the total number of sessions. This is introduced to penalize topics that are commonly used, e.g. topic 20. The idea is to give less score to these topics since they are used frequently anyway while other less prominent topics can be captured in the model. The feature vectors are then normalized before being fed into the K-Means clustering algorithm. We used the Weka machine learning library for the clustering task [9]. The generated clusters correspond to the underlying usage states.

They have distinctive topic distributions and serves as estimate of the final usage states of the users.

## DISCUSSION OF FINDINGS

### Apps and Sessions

The top five most used apps are SMS, Phone, Mail, Facebook, and Safari. Apps have mean usage duration of 64.85 seconds (SD = 527.99), which is considered normal for apps usage behavior. Facebook and Safari tend to be used for longer duration, with mean duration above 1 minute, while apps like SMS, Phone, and Mail all have mean duration below 1 minute. This is unsurprising as users spend more time consuming contents available on Facebook and web pages.

Overall, we identified 273,912 usage sessions with an average of 2.12 apps per session (SD = 2.06). This shows that users tend to use small number of apps per session, which pertains to particular utility. The overall average duration of usage session is 172.8 seconds (SD = 730.2), with the minimum duration of several seconds and the maximum of 11 hours (Alarm clock apps). This shows a strong diversity in usage behavior. High variability in usage time from all 24 users suggests that personalization is needed on a user level.

We also explored the apps usage patterns in the sessions. We found a total of 38,145 unique patterns. The most frequent apps usage pattern is a consecutive sequence of SMS apps, which tops for sessions of sizes $1 - 4$. The top 100 apps patterns have 1.9 apps on average (SD = 0.77), which shows that usage sessions usually include a small number of apps.

We examined K-Unique apps patterns in the usage trace. For example, $S_1$ = (SMS, SMS, SMS) is a 1-Unique apps pattern and $S_2$ = (SMS, Facebook, Mail) is a 3-Unique apps pattern. We found that 17,136 (95%) of all unique patterns consisted of at least two unique apps and only 784 are 1-Unique app patterns. 3-Unique apps patterns have the highest occurrence and account for 27% of total K-Unique apps patterns. This indicates co-occurrence patterns in apps usage behavior.

In the usage sessions topic 20 *communication* is the most prominent (60% of all sessions) with 164,352 counts, followed by topic 16 with 60,740 usages (22 %). The next most frequent topics are working (16%) and social topics (15%).

### Usage State

We expect users to engage in multiple topics of apps within a usage session. The K-Means clustering algorithm produces a distribution of types of usage sessions based on app topics. We expect each cluster to capitalize on one or more topics to represent a usage behavior. In the following we analyze the usage behavior in terms of usage states first across users and then for each user separately.

#### *Between-Users Analysis*

We first combined the usage traces of all users to find the corresponding representative clusters, which correspond to the usage states. In total, there were 13 distinct usage clusters across users as shown in Table 2. We analyzed the centroids of each cluster to exploit its characteristics. Cluster centroids are simply the means of each instance within the cluster. From the centroids we observed that there are one or more prominent topics that make up the cluster, e.g. Cluster S5 reflects tasks such as social networking. The top 3 topics alone accounts for between 47% - 96% of the total topic distribution within the clusters. This means that generally, mobile usage states cover a small set of focused topics. Communication is present in all usage states, which means that users are likely to use apps like SMS and Email in all states.

**Coherence Validation**: We measure the coherence of the usage states generated by our approach by manually analyzing the topic distribution for each centroid. We checked whether the co-occurrences of topics make sense that is whether the topics and their distributions would reflect a meaningful usage scenario. From the centroid of each usage state, we extracted a set of topics, whose distributions are greater than 0.10. We considered topics below this threshold to be negligible. We then asked 4 smartphone users to manually rate the coherence of the usage states on a 5-level scale (from very good to very bad). We then converted the individual ratings for each usage state into a numerical scale ([-2, 2] range). 46% of usage states fall between 1-2 scale, while another 46% fall between 0-1 scale. The usage states have a $good$ to $neutral$ coherence, with an exception of state S8, which have average coherence score of -0.5. This is because state S8 contains 3 topics (Topic 27, 11, & 9) where the users found to be slightly contradicting to each other.

**Long stretches**: We also examined long stretches, i.e. day-sequences of usage states. On average, users have 34.5 usage states (SD = 35.49) per day. The number of unique states per day is only 5.88 (SD = 2.83). Users tend to start using their phone around 7-8 AM in a communication state and end at 12 AM-1 AM with device feature state, e.g. using preference or timer app. Large standard deviation value of usage state shows high diversity in usage behavior therefore researcher should adapt to individual user behaviors. Most usage states are used between afternoon and evening hours, with an exception for state S11 which is more prominent during morning and late night hours. The top occurring unique state pattern accounts for only 2% for all day-sequences (sequences of S0 and S11). This shows that across users, there are no definite long stretch patterns in usage states on daily basis.

**Routines**: To understand routines that exist in the data, we extracted 1-hour sequences (chain of usage states of 1 hour long). 1-hour sequences have 3.9 usage states on average (SD = 2.72) for weekdays and 3.75 usage states (SD = 2.63) for weekends, which is expected as users uses their phone 2-4 times an hour. Most sequences are only of 1 single usage state. For multi-usage state sequences, the most notable unique-state patterns were sequences of combination of communication activities with gaming, social, and device feature states, e.g. S0-S5 and S0-S8 for morning and evening hours. Browsing (S3) and Social (S5) activities are more prominent during afternoon hours (12 PM-16 PM).

We then examine the most occurring start-end state patterns. Morning routines are mostly marked by communication and gaming activities; sequences usually start and end with communication and gaming state. Weekday mornings have an

| Usage State | Top Topics | Top Apps |
|---|---|---|
| S0 | Communicating | SMS |
| S1 | Shopping | Groupon |
| S2 | Relaxing | Media Player |
| S3 | Browsing,Search | Safari |
| S4 | Phone-related | MobilePhone |
| S5 | Social | Facebook |
| S6 | Finance | SMS |
| S7 | iOS Utilities | Calender |
| S8 | Gaming, Support, Weather | WordsWithFriends, Weather |
| S9 | Navigating | Maps |
| S10 | Photography, Task-Planning | Camera, Notes |
| S11 | Device Feature | Timer |
| S12 | Communicating, Phone-related | SMS,Email,Phone |

Table 2. **Top topics/apps per cluster for all users $minSup$=0.05.**

additional prominent state of S11 (timer app) which is less apparent on weekend. Afternoon and evening routines are mostly related to communication and gaming activities.

### *Within-Users Analysis*

We examine the output clusters based on individual usage behavior (user-specific pause time) with the same LDA and $minSup$ parameters. 24 users have on average 13.7 usage states (SD=4.17, min=5, max=25). Although the mean value is close to the number of usage states across users, the 95% confidence interval is 12 to 15.4 usage states. This means that it is still better if we study fine-grained behavior for specific users rather than using a global usage state model. We also calculated the average Jaccard Index [17] between each user's specific model with the global model based on the centroids. We consider two centroids to be equal if their distance is less than a threshold. The Jaccard Index is 0.28 (SD = 0.065), with 95% confidence interval of 0.24-0.29. This means that in terms of their centroids the usage models of each user vary much from the global model.

Although, the number of usage states may vary across users, we find that there are two common usage states across users: S0 and S4. We find this by assigning cluster centroids into groups. We limit that each centroid in a group cannot be further than a certain threshold to other members within the group. S8 (gaming usage state) is found common across 17 users, while S2 (relaxing state) is found across 12 users.

## RELATED WORK

User behavior on mobile devices has been shown to exhibit some regular temporal, spatial, and usage pattern [20, 19, 3, 6, 14, 5] and the usage behavior of mobile devices changes under different contexts [4, 21, 16]. The usage of apps exhibit specific temporal and geographical pattern [20, 3], e.g., SMS and Phone are shown to have an evenly distributed pattern [19], while apps like news or weather apps are used more frequently in morning hours, and music or video player apps are used more frequently on the move.

Xu et al. [20] profiled usage patterns of mobile apps based on apps category. They find that some apps have high probability of usage co-occurrence and that users tend use several alternative apps within the same category. Therefore we argue that it is better to look at apps behavior from topics level to understand the usage functionality. Falaki et al. [7] showed that a strong diversity exists in usage behavior, e.g. the range of interaction time lasts from 30-500 minutes a day and consists of 10-200 apps sessions per day, while each session can last from several minutes to an hour. The authors emphasized that personalization is needed on a user level not group level because of the strong diversity in usage behavior. Our results confirm this result as we showed that user specific model would better represent usage states. Furthermore, based on the study of usage behavior from 4,100 android users, Böhmer et al. [3] showed that communication related apps are used first when user turn on their phone. The usage transition probability to communication apps is also high, which is similar to our finding that communication tasks are the most prominent in mobile usage behavior.

However, these studies did not attempt to characterize the notion of usage "states" that exists in the mobile usage traces, which is the focus of our research. Previous researches have used LDA on app description [10] and app reviews [8] to extract app features. Other researchers also tried to automatically categorize the apps using LDA on app pages [18]. Our work also proposed to use LDA to generate topics based on app descriptions. We then combine these topics with the actual usage traces of the app to derive usage states, which characterizes recurring usage situations. We studied the mobile usage behavior focusing on more general usage states and topics rather than concrete apps or their specific categories. We investigated the variability of usage behavior and illustrated that usage states exist between usage sessions. Our results give first insights that these usage states can represent the user's behavior on mobile devices.

## CONCLUSION

We introduced a method for identifying user behavioral states from mobile usage traces. We used LDA Topic Modeling, Frequent Itemset mining, and K-means Clustering algorithms in order to generate the behavioral states of a user. The states are repetitive, meaningful, and have interesting characteristics. Identifying and predicting them could be beneficial for personalized predictive and recommendation services. Our work gives first preliminary evidence that mobile usage behavior can be segmented into meaningful repeatable states.

Our approach and study results have several limitations. App descriptions can be ambiguous or even misleading. Furthermore, the results are strongly dependent on the Livelab dataset, which, e.g., did not contain apps for categories like Catalogs, Food& Drink, and Newsstand at the time of its collection. Future work should also re-adjust the number of topics to account for additional categories within the dataset. Other limitations include the parameters settings such as the $minSup$ and LDA threshold. Different values for these parameters can affect the outcome of the final usage states.

## ACKNOWLEDGEMENTS

We thank the reviewers for the constructive reviews. This work is part of the EU research project MUSES FP7-318508.

## REFERENCES


1. Rakesh Agrawal and Ramakrishnan Srikant. 1994. Fast Algorithms for Mining Association Rules in Large Databases. In *Proceedings of 20th International Conference on Very Large Data Bases (VLDB'94)*. 487–499.

2. David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022. **http://www.jmlr.org/papers/v3/blei03a.html**

3. Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger, and Gernot Bauer. 2011. Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage. In *Proceedings of the 13th Conference on Human-Computer Interaction with Mobile Devices and Services*. 47–56.

4. Trinh Minh Tri Do, Jan Blom, and Daniel Gatica-Perez. 2011. Smartphone Usage in the Wild: A Large-scale Analysis of Applications and Context. In *Proceedings of the 13th International Conference on Multimodal Interfaces (ICMI '11)*. ACM, 353–360. DOI: **http://dx.doi.org/10.1145/2070481.2070550**

5. Trinh-Minh-Tri Do and Daniel Gatica-Perez. 2010. By Their Apps You Shall Understand Them: Mining Large-scale Patterns of Mobile Phone Usage. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia (MUM '10)*. ACM, Article 27, 10 pages.

6. Nathan Eagle and AlexSandy Pentland. 2009. Eigenbehaviors: identifying structure in routine. *Behavioral Ecology and Sociobiology* 63, 7 (2009), 1057–1066. DOI: **http://dx.doi.org/10.1007/s00265-009-0739-0**

7. Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. Diversity in Smartphone Usage. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys '10)*. ACM, 179–194.

8. E. Guzman and W. Maalej. 2014. How Do Users Like This Feature? A Fine Grained Sentiment Analysis of App Reviews. In *Requirements Engineering Conference (RE), 2014 IEEE 22nd International*. 153–162. DOI: **http://dx.doi.org/10.1109/RE.2014.6912257**

9. Mark A. Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA data mining software: an update. *SIGKDD Explorations* 11, 1 (2009), 10–18. DOI: **http://dx.doi.org/10.1145/1656274.1656278**

10. M. Harman, Yue Jia, and Yuanyuan Zhang. 2012. App store mining and analysis: MSR for app stores. In *Proc. of Working Conference on Mining Software Repositories - MSR '12*. 108–111. DOI: **http://dx.doi.org/10.1109/MSR.2012.6224306**

11. Walid Maalej, Hans-Jörg Happel, and Asarnusch Rashid. 2009. When users become collaborators: towards continuous and context-aware user input. In *Proceeding of the 24th ACM SIGPLAN conference companion on Object oriented programming systems languages and applications (OOPSLA '09)*. ACM, 981–990. DOI: **http://dx.doi.org/10.1145/1639950.1640068**

12. Walid Maalej and Dennis Pagano. 2011. On the Socialness of Software. In *Proceedings of the International Software on Social Computing and its Applications* (2011-01-01). IEEE Computer Society.

13. Andrew Kachites McCallum. 2002. MALLET: A Machine Learning for Language Toolkit. **http://mallet.cs.umass.edu**

14. Abhinav Parate, Matthias Böhmer, David Chu, Deepak Ganesan, and Benjamin M. Marlin. 2013. Practical Prediction and Prefetch for Faster Access to Applications on Mobile Phones. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, 275–284. DOI: **http://dx.doi.org/10.1145/2493432.2493490**

15. Clayton Shepard, Ahmad Rahmati, Chad Tossell, Lin Zhong, and Philip T. Kortum. 2010. LiveLab: measuring wireless networks and smartphone users in the field. *SIGMETRICS Performance Evaluation Review* 38, 3 (2010), 15–20.

16. Choonsung Shin, Jin-Hyuk Hong, and Anind K. Dey. 2012. Understanding and Prediction of Mobile Application Usage for Smart Phones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, 173–182. DOI: **http://dx.doi.org/10.1145/2370216.2370243**

17. Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. 2005. *Introduction to Data Mining*. Addison-Wesley.

18. Svitlana Vakulenko, Oliver Müller, and Jan vom Brocke. 2014. Enriching iTunes App Store Categories via Topic Modeling. In *Proceedings of the 35th International Conference on Information Systems*.

19. Hannu Verkasalo. 2009. Contextual patterns in mobile service usage. *Personal and Ubiquitous Computing* 13, 5 (2009), 331–342.

20. Qiang Xu, Jeffrey Erman, Alexandre Gerber, Zhuoqing Morley Mao, Jeffrey Pang, and Shobha Venkataraman. 2011. Identifying diverse usage behaviors of smartphone apps. In *Internet Measurement Conference*. 329–344.

21. Tingxin Yan, David Chu, Deepak Ganesan, Aman Kansal, and Jie Liu. 2012. Fast App Launching for Mobile Devices Using Predictive User Context. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys '12)*. ACM, 113–126.