The Need to Account for Geographical Diversities in Mobile Data Research

Akhil Mathur
Bell Labs, Nokia
akhil@bell-labs.com

Fahim Kawsar
Bell Labs, Nokia
fahim.kawsar@bell-labs.com

ABSTRACT

Large-scale mobile data studies have the potential to provide valuable insights regarding usage behavior of smartphone users. A major challenge in generalising the findings of these studies is the inherent population diversity in large-scale smartphone usage data. Many published studies however do not account for population heterogeneities in their analysis, and instead present their findings based on the aggregate usage data. In this paper, we investigate the effects of geographical diversity on smartphone data collected from 130 users from India, Europe and the US over a period of four months. Our results show significant differences in daily usage duration, session-level usage, and temporal usage patterns across various geographies, and stress the need to account for population heterogeneities in mobile data research and its application in real-world systems.

Keywords

Mobile Data, Population Diversity, User Study, Android

1. INTRODUCTION

Mobile phones have evolved from simple communication tools into powerful information, communication and entertainment devices. By the year 2020, it is projected that 5.4 billion people in the world will have a mobile phone, higher than those projected to have electricity (5.3 billion), running water (3.5 billion) or cars (2.8 billion) [2]. Across different app stores, there are more than 4 million applications that the users can download on their phones [1], many of them capable of capturing rich data about users and their environment. By accurately inferring user and environment context from this data, these apps can provide personalised, contextual services to the end-users and substantially improve the user experience.

Consequently, we have seen a number of studies that have collected smartphone usage data at scale, and analysed it to infer usage patterns. Falaki et al. [8] conducted a comprehensive study of smartphone usage to characterise user interactions with the device and various applications, and the impact of these interactions on network and energy consumption. Bohmer et al. [6] extended their analysis to understand contextual factors that influence application usage – for example, they found that "news applications are most popular in the morning and games are at night, but communication applications dominate through most of the day". Other works have studied the effect of mobility on usage [21], notification delivery and reception on smartphones [15, 17, 23], energy consumption [12], and predicting application usage [25].

A major challenge in generalising the findings of these user studies is the inherent population diversity in large-scale smartphone usage data. In the human-computer interaction community, researchers aim to control for population diversity in their experiment design, e.g., by selecting participants with similar skills or demographics. In mobile data research however, this issue of population diversity is even more critical – as many of the large-scale data collection studies are done by publishing an application on app stores (e.g. Apple Store or Google Play) [6, 15, 23], it indeed becomes difficult to control for user diversity in the data.

Our aim, in this work, is to highlight one particular source of user diversity (viz. geographical diversity) – and how it impacts smartphone usage. While it is not surprising to expect diversity in smartphone usage due to geographical variations, we observe that many published studies on mobile data (e.g. [6, 23, 25]) do not account for such variations even when analysing the data. Other studies have acknowledged the user diversity (e.g. [8]), however they did not provide in-depth analysis of how it affects smartphone usage.

In this paper – by analysing the data gathered in a 130-user longitudinal study, we demonstrate several key differences in smartphone usage between users in India, Europe and United States. We also contrast our findings with past research on smartphone usage in another country (i.e., Korea). We found that there are significant differences across various geographies in terms of daily usage duration, session-level usage, and temporal usage patterns – and if these differences are left unaccounted for, they may result in very misleading findings about real-world smartphone usage. The paper concludes by discussing implications of these geographical variations for mobile data studies, application development, and mobile-based data inference models.

MobiData'16, June 30 2016, Singapore, Singapore

ISBN 978-1-4503-4327-5/16/06. . . $15.00

DOI: `http://dx.doi.org/10.1145/2935755.2935761`

2. RELATED WORK

In this section, we overview two categories of past research that are relevant to our work on investigating population heterogeneities in mobile data.

Mobile Data Studies. Over the last few years, there have been a number of studies on mobile data usage in-the-wild. Falaki et al. [8] conducted a comprehensive study of smartphone usage to characterise user interactions with the phone, application usage, network traffic and energy usage. Bohmer et al. [6] analysed contextual factors that influence application usage. Oulasvirta et al. [16] discussed the possibility of mobile users developing a checking habit that involves brief and frequent content consumption.

Various techniques have been proposed [19, 20, 22] for continuously and unobtrusively gathering mobile usage and sensor data. Based on such data, researchers have proposed smart home screens and launcher apps to improve user experience [4, 28] and device battery life [9], as well as developed inference models for predicting the user's next application [25, 26]. Other works have studied the effect of mobility on usage [21], notification delivery and reception on smartphones [15, 17, 23], and energy consumption [12].

Data Heterogeneities. Researchers have identified various hardware and software heterogeneities that impact inference techniques on mobile phones. Blunck et al. [5] discussed how variations in GPS duty cycling across platforms can adversely affect data quality and performance of inference models. Stisen et al. [27] highlighted that run-time factors such as instantaneous I/O load can lead to an unstable sensor sampling rate, thereby compromising the data quality. Das et al. [7] found significant hardware variations in acoustic components such as microphones and micro-speakers of a mobile device.

As a solution to such device heterogeneities, Lu et al. [13] proposed a framework which can identify stress using microphones in diverse acoustic environments. Lu et al. [14] presented a continuos sensing platform robust to accelerometer variations and biases. To account for population diversity in sensor data, Lane et al. [10] looked at incorporating inter-person similarity measurements from crowd-sourced sensor-data into the classifier training process.

3. RESEARCH STUDY

In this section, we present our data collection system, and give an overview of the participant demographics and our data analysis methodology.

3.1 System

In order to collect smartphone usage data, we developed an Android application and released it on Google Play. The app, implemented for Android 5.0+, runs as a background service on the user's phone and passively records all usage sessions on the device. More specifically, the application collects data on *Screen Events* (i.e., times when the screen was turned on, off, or unlocked), *App Events* (i.e., times when an app is in foreground or background), *Call Events* (i.e., times of incoming, outgoing, or missed calls) and *Notification Events* (times when a notification is received, read, or dismissed).

When the app is first launched, users are asked to fill a short demographic questionnaire in which they input their age, gender, and country. Next, users are guided through a permission wizard wherein they provide permissions to the app for accessing notifications and application usage data on their phone. Both these permissions are categorised as *system-level permissions* in Android, and and must be manually approved by the user from the Settings Menu on their phone. The apps guides the user through this approval process, and once the permissions are granted – the app starts running in the background to collect data logs. The logs are stored locally on the external storage of the phone, and are periodically uploaded to a remote server.

3.2 Participants and Data

After publishing the app on Google Play in December 2015, we advertised it on social networks and email lists to solicit participation. In total, the app was installed by 130 users from 12 different countries. The raw data set obtained from the users until April 2, 2016 contains 40,414 call events, 630,241 notification events, 890,945 application events, and 901,624 screen events.

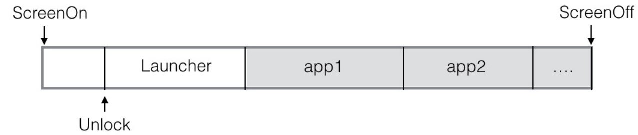

Figure 1: Example of a smartphone usage session.

Out of 130 users, we selected 86 users (25 female) for our analysis who were i) based in our target geographies, i.e. India, Europe or the US, and ii) had kept the app installed on their phone for at least one month, which we believe is a sufficient duration to get representative usage data.

The shortlisted users were mostly young adults aged between 18 to 30 years, and 65% were students. Of the 86 users, 54 were from India, and the remaining 32 were from 6 different European countries and the US. For our geographical analysis, instead of analysing data at a per-country level, we decided to group the users from Europe and US into one category called *'Western Users'* and compared their usage to *'Indian Users'*. There were primarily two reasons for this choice: i) to ensure that we have sufficient number of users in each comparison group, and ii) Europe and US, although geographically separate, are both developed regions with high smartphone penetration for a number of years. On the contrary, smartphone adoption in India is a relatively recent phenomenon, and the smartphone penetration is much lower [3]. Hence, even with the aforementioned categorisation, it will be possible to contrast the smartphone usage among two diverse groups.

3.3 Data Analysis

Similar to the prior work on mobile usage analysis, we also centred our analysis around a usage session as shown in Figure 1. A usage session is defined as the period between screen-on and screen-off – as such, we processed the *Screen Events* in the logs to generate usage sessions. On Android, events such as notification arrival and battery charging can switch on the screen automatically – therefore, for an accurate representation of usage sessions, we filtered out the sessions that did not elicit any user interactions (i.e., screen unlocking). In each session, a user can access 0, 1 or multiple applications – as such, we processed the *Application Events* in the logs to find the apps used in each session.

Smartphone usage sessions can be triggered by external cues (e.g., through a notification or a call), or internal cues (i.e., user motivated). To differentiate the two kinds of sessions, the *notification* and *screen* events were processed together; we categorised a session as 'externally cued' if i) the session was initiated within 30 seconds of any notification arrival, or ii) if any notifications from the apps used in the session arrived between the current session and the preceding session [11]. All other sessions were categorised as 'internally cued'.

Finally, before comparing the usage between Indian and Western users, we measured the skewness of each comparison metric (e.g. session duration, session count, application count) within the group, and found that the data distributions for all metrics were within the bounds of a normal distribution ($-1 < skewness < 1$).

4. STUDY RESULTS

We now analyse the aggregated usage, session-level usage, and temporal usage across the two user groups (Indian users and Western users), and also contrast them with prior research on mobile data usage.

4.1 Aggregate Usage

We begin by comparing smartphone usage of the two groups in our dataset (Indian and Western) through unpaired two-tailed t-test, and report the p value, t statistic and $Cohen's$ d. In addition, we also contrast the usage of these two groups with the findings of Lee et al. [11] regarding smartphone usage of young 'Korean' users. In their study, Lee et al. [11] investigated smartphone addiction among Korean students; as such, they collected smartphone usage data from 'addicted users' and 'regular users' for nearly one month. For our geographical analysis, we use the data from 'regular users' in [11] and compare it to the smartphone usage of participants in our dataset. The participants in our dataset also belong to a similar age group as those in [11] and majority of them are students; therefore, we believe it is reasonable to compare the two datasets.

	Indian	Western	p	t	d	Korean [11]
Usage duration (mins) per day	149	113	0.04*	0.51	1.96	207
Sessions per day	91	66	0.01*	2.6	0.56	100.1
Unique apps	80.5	82.4	0.8	-0.15	0.1	NA
Entropy top-5	3.42	3.55	0.38	-0.87	0.18	1.96
Entropy top-10	3.89	4.03	0.34	-0.95	0.2	2.53

Table 1: Comparison of aggregate usage across user groups. Results of unpaired two-tailed t-test are provided for groups in our dataset (* indicates significant difference between groups, $p < 0.05$)

Table 1 shows the aggregate usage across different groups. We found a significant difference ($t = 1.96$, $p < 0.05$, $Cohen's$ d = 0.51) between the mean usage duration in a day between Indian users and Western users. While Indian users had a mean usage duration of 149 minutes per day, western participants used the phone for 113 minutes per day on an average. In contrast, the Korean users showed a much higher mean usage duration of 207 minutes per day. Similarly, we found a significant difference between the number of usage sessions per day across Indian and western users, with Indian users having nearly 38% more usage sessions in a day. Korean users from [11] had higher usage sessions than both the groups in our dataset. There was no significant difference in the number of unique apps used by both groups during the study ($p > 0.05$).

Next, we investigate the diversity in the usage of top-k apps (ordered by usage duration) across users. We compute the Shannon Entropy as a measure of *Evenness* of application usage. Shannon Entropy for top-k apps denoted by H' is calculated using the equation below, where d_i denotes the relative usage duration of i^{th} app.

$$H' = -\sum_{i=1}^{k} d_i.log(d_i)$$

A higher value of H' indicates that usage was spread evenly across all applications, whereas a lower value of H' indicates that usage was skewed towards a small number of applications.

We found no significant difference in usage diversity in top-5 or top-10 apps between Indian and Western user groups. However, the Shannon Entropy value reported for the Korean users was lower than our findings, suggesting that Korean users have more skewed usage patterns for top-5 and top-10 apps.

4.2 Session-level Usage

Table 2 shows the comparison of session-level usage across different groups. Our analysis did not reveal any significant difference in the mean session duration between Indian and Western users ($p = 0.8$), however we observe that Korean users have 30% higher mean session duration than the other two groups. We did observe a significant difference in the inter-session gap between the two groups ($p < 0.02$); results showing that Indian users start a usage session every 9.6 minutes on average as compared to 16.6 minutes for Western users and 13.61 minutes for Korean users.

	Indian	Western	p	t	d	Korean [11]
Mean session duration (s)	97.8	95.3	0.8	0.22	0.05	129.9
Inter-session gap (min)	9.6	16.6	0.02*	2.78	1.07	13.61
App launches per session	3	2.4	0.01*	2.95	0.58	3.16
Unique apps per session	1.91	1.71	0.01*	2.71	0.55	NA

Table 2: Comparison of session-level usage across groups. Results of unpaired two-tailed t-test are provided for groups in our dataset (* indicates significant difference between groups, $p < 0.05$)

Next, we analysed the app launch count in each session (i.e. how many apps were launched in a session, including repetitions). A significant difference was found ($p < 0.005$), and results show that Indian users launch more apps (mean = 3.0) per session than Western users (mean = 2.4). Similarly, we found a significant difference ($p < 0.01$) in number of unique apps used in a session, with Indian users exhibiting a higher unique app count.

4.3 Session Analysis

Now we analyse the nature of individual sessions. As we described in the previous section, we categorised the sessions into two classes: internal and external depending on how they were initiated. Both user groups had a similar distribution of sessions, with external sessions accounting for 17.7% of all sessions for Indian users and 18.1% of sessions for Western users. Lee et al. [11] however reported a very different session distribution for Korean users, where externally cued sessions accounted for 79% of all sessions!

In order to account for varying session durations, we further categorised the sessions into three classes (following the methodology of Banovic et al. [4]): glance, short, long. *Glance* sessions are those sessions in which no apps are used – here a user may unlock the phone to glance at a notification (or check the time), but does not launch any application. A *Short* session is less than 60 seconds in duration, and involves the use of 1 or 2 apps, whereas a *Long* session lasts for more than 60 seconds and may involve 1 or multiple apps. Therefore in total, we had six categories for sessions (3 based on duration and usage x 2 based on initiation type).

	Internal		**External**	
	Indian	Western	Indian	Western
Glance	53.4%	45.9%	8%*	5.2%*
Short	11.1%*	16.4%*	4%	5.3%
Long	15.9%	19%	5.5%	6.2%

Table 3: Comparison of session types across groups. Cells marked with * indicates a significant difference between groups ($p < 0.05$).

Table 3 presents the session distribution across both groups. We found a significant difference in the percentage of External-Glance ($t = 2.27$, $d = 0.62$, $p < 0.05$) sessions between the two groups. This suggests that on receiving notifications (i.e., external cues), Indian users show a higher preference to glance at its content, rather than opening the application associated with it. We also found a significant difference in the percentage of Internal-Short session ($t = -2.6$, $d = 0.76$, $p = 0.01$) between the two groups, suggesting that Western users self-initiate more short duration sessions than Indians.

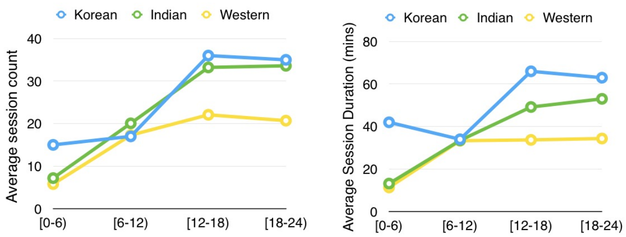

Figure 2: Left: Temporal view of session counts across groups. Right: Temporal view of session durations across groups.

4.4 Temporal Analysis

Finally, we analyse the temporal usage patterns in a day for all the groups. To facilitate comparison with Lee et al. [11], we divided the day in four equal buckets of six hours, and grouped the smartphone usage by buckets. Figure 2 presents the temporal distribution of session counts and session durations across the user groups.

As evident from the figure 2 , we found significant differences both in session duration and session count between 12noon - 6pm and 6pm - 12am (all $p < 0.05$), with Indian users having higher session counts and durations at these times than Western users. Korean users exhibited higher session duration and session count than the other groups during 12noon - 6pm and 6pm - 12am.

5. DISCUSSION

We now discuss the implications of our findings for mobile data studies, and also broadly highlight the challenges for application design and user modelling due to geographical variations in mobile data.

Implications for Mobile Data Research. We present two empirical examples which clearly highlight the challenges posed by geographical diversity for mobile data research. Firstly, we found that the percentage of externally cued sessions for Korean users (79%) was much higher than for Indian (17.7%) and Western users (18.1%). This clearly suggests that app developers should focus on designing better external cues (i.e., notifications) to increase smartphone engagement among Korean users. Now hypothetically, if this analysis on session initiation was done over aggregate data from the three user groups (i.e. without accounting for geographical diversities) – it would have presented a misleading picture on session initiation behavior of the users, and may not have resulted in a concrete design guideline for app developers.

In another example, we look at the large-scale smartphone data collection exercise by Bohmer et al. [6] which attracted participation of users from 20+ countries. In their analysis on the aggregate data, they found that almost half (49.8%) of all smartphone sessions were shorter than 5 seconds. On the other hand, our dataset reveals that the percentage of micro-sessions (i.e., sessions shorter than 5 seconds) was only around 22% for both Indian and Western users (with no significant difference between the two groups). The reason behind this very contrasting result however remains unclear: were there some countries in the dataset of [6], where users have particularly short attention spans, which in turn affected the aggregate statistics presented in their paper? Indeed, a geographical analysis of the data would have shed more light on these usage variations between the datasets.

This contrast in findings is also likely to confuse application or platform developers: should they design the mobile software platform to account for very short attention spans of the user (as reported by [6]), e.g. by creating and displaying content summaries on the lock screen? Or is the current design of mobile notifications – where minimal information about the content is shown on the lock screen – a better way to entice the users to engage with the apps?

Implications for User Modeling. The geographical variations in smartphone usage can also impact the accuracy of machine learning inference models in-the-wild. While many works focus on building personalised inference models (e.g., [25, 28]), it is also common to develop models from composite data (e.g., [18]), in order to avoid the user cold-start problem [24]. For such composite models, population heterogeneities (e.g., geography, age) could be a major obstacle in their wider applicability. For example, Lee et al. [11] developed a classification model to infer smartphone addiction based on smartphone usage statistics of Korean users. If the same model has to be applied to Indian and Western users – who have significantly different usage behavior than Korean users as shown in our study – it would require substantial tuning of the model parameters to account for the geographical variations.

Tackling Other Sources of Heterogeneities. In this paper, we focused on one particular source of diversity in mobile data (viz. geographical) - however there are multiple other sources of heterogeneities that should be looked into in detail. For example, gender, age, or profession of a user is likely to have a significant impact on their usage behavior. However, most studies on mobile data usage ignore these heterogeneities while presenting aggregate statistics or building inference models. We can certainly take inspiration from the mobile sensing literature, wherein researchers have developed techniques to account for hardware [13, 27], software and orientation based heterogeneities [14, 27] in on-device sensors.

6. CONCLUSIONS

The immense growth in the smartphone industry, proliferation of mobile app stores, and availability of software frameworks for continuously gathering mobile usage and sensor data has enabled us to study mobile device usage at an unprecedented scale. However, for mobile app and platform developers to confidently apply these research findings to develop real-world systems, it is extremely important to improve the external validity and generalisability of mobile data studies. In this paper, we showed that geographical diversity in the population can lead to significant differences in smartphone usage patterns, which in turn, can affect the external validity of mobile data studies and even lead to misleading results about user behavior. As such, we urge fellow mobile data researchers to account for population heterogeneities while analysing smartphone usage data or developing mobile-based data inference models. We also hope that our findings on geographical variations in mobile data can trigger interesting discussions at the workshop, which lead to design of better mobile data studies and systems that are robust against population variations.

7. REFERENCES

[1] App store statistics. `http://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/`. Accessed: 2016-04-10.

[2] Cisco smartphone prediction. `https://newsroom.cisco.com/press-release-content?articleId=1741352`. Accessed: 2016-04-10.

[3] Smartphone penetration in India. `http://www.statista.com/statistics/257048/smartphone-user-penetration-in-india/`. Accessed: 2016-04-10.

[4] N. Banovic, C. Brant, J. Mankoff, and A. Dey. Proactivetasks: the short of mobile device use sessions. In *Proceedings of MobileHCI '14*, pages 243–252. ACM, 2014.

[5] H. Blunck, N. O. Bouvin, T. Franke, K. Grønbæk, M. B. Kjaergaard, P. Lukowicz, and M. Wüstenberg. On heterogeneity in mobile sensing applications aiming at representative data collection. In *Proceedings of UbiComp '13 Adjunct*, pages 1087–1098. ACM, 2013.

[6] M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer. Falling asleep with angry birds, facebook and kindle: a large scale study on mobile application usage. In *Proceedings of MobileHCI '11*, pages 47–56. ACM, 2011.

[7] A. Das, N. Borisov, and M. Caesar. Fingerprinting smart devices through embedded acoustic components. *arXiv preprint arXiv:1403.3366*, 2014.

[8] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in smartphone usage. In *Proceedings of MobiSys '10*, pages 179–194. ACM, 2010.

[9] D. Ferreira, E. Ferreira, J. Goncalves, V. Kostakos, and A. K. Dey. Revisiting human-battery interaction with an interactive battery interface. In *Proceedings of UbiComp '13*, pages 563–572. ACM, 2013.

[10] N. D. Lane, Y. Xu, H. Lu, S. Hu, T. Choudhury, A. T. Campbell, and F. Zhao. Community similarity networks. *Personal Ubiquitous Comput.*, 18(2):355–368, Feb. 2014.

[11] U. Lee, J. Lee, M. Ko, C. Lee, Y. Kim, S. Yang, K. Yatani, G. Gweon, K.-M. Chung, and J. Song. Hooked on smartphones: an exploratory study on smartphone overuse among college students. In *Proceedings of CHI '14*, pages 2327–2336. ACM, 2014.

[12] D. Li, S. Hao, J. Gui, and W. G. Halfond. An empirical study of the energy consumption of android applications. In *2014 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 121–130. IEEE, 2014.

[13] H. Lu, D. Frauendorfer, M. Rabbi, M. S. Mast, G. T. Chittaranjan, A. T. Campbell, D. Gatica-Perez, and T. Choudhury. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of UbiComp '12*, pages 351–360. ACM, 2012.

[14] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell. The jigsaw continuous sensing engine for mobile phone applications. In *Proceedings of the 8th ACM conference on embedded networked sensor systems*, pages 71–84. ACM, 2010.

[15] A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley, and M. Musolesi. My phone and me: Understanding people's receptivity to mobile notifications. 2016.

[16] A. Oulasvirta, T. Rattenbury, L. Ma, and E. Raita. Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing*, 16(1):105–114, 2012.

[17] V. Pejovic and M. Musolesi. Interruptme: designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of UbiComp '14*, pages 897–908. ACM, 2014.

[18] M. Pielot, T. Dingler, J. S. Pedro, and N. Oliver. When attention is not scarce - detecting boredom from mobile phone usage. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '15, pages 825–836, New York, NY, USA, 2015. ACM.

[19] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas. Emotionsense: a mobile phones based adaptive platform for experimental social psychology research. In *Proceedings of UbiComp '10*, pages 281–290. ACM, 2010.

[20] M. Raento, A. Oulasvirta, R. Petit, and H. Toivonen. Contextphone: A prototyping platform for context-aware mobile applications. *Pervasive Computing, IEEE*, 4(2):51–59, 2005.

[21] A. Rahmati and L. Zhong. Studying smartphone usage: Lessons from a four-month field study. *Mobile Computing, IEEE Transactions on*, 12(7):1417–1427, 2013.

[22] N. Ramanathan, F. Alquaddoomi, H. Falaki, D. George, C.-K. Hsieh, J. Jenkins, C. Ketcham, B. Longstaff, J. Ooms, J. Selsky, et al. Ohmage: an open mobile system for activity and experience sampling. In *PervasiveHealth*, pages 203–204. IEEE, 2012.

[23] A. Sahami Shirazi, N. Henze, T. Dingler, M. Pielot, D. Weber, and A. Schmidt. Large-scale assessment of mobile notifications. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 3055–3064. ACM, 2014.

[24] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock. Methods and metrics for cold-start recommendations. In *Proceedings of SIGIR '02*, pages 253–260. ACM, 2002.

[25] C. Shin, J.-H. Hong, and A. K. Dey. Understanding and prediction of mobile application usage for smart phones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 173–182. ACM, 2012.

[26] V. Srinivasan, S. Moghaddam, A. Mukherji, K. K. Rachuri, C. Xu, and E. M. Tapia. Mobileminer: Mining your frequent patterns on your phone. In *Proceedings of UbiComp '14*, pages 389–400. ACM, 2014.

[27] A. Stisen, H. Blunck, S. Bhattacharya, T. S. Prentow, M. B. Kjærgaard, A. Dey, T. Sonne, and M. M. Jensen. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 127–140. ACM, 2015.

[28] T. Yan, D. Chu, D. Ganesan, A. Kansal, and J. Liu. Fast app launching for mobile devices using predictive user context. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 113–126. ACM, 2012.