

Describing Patterns and Disruptions in Large Scale Mobile App Usage Data

Steven Van Canneyt
Ghent University
steven.vancanneyt@ugent.be

Marc Bron
Yahoo
mbron@yahoo-inc.com

Andy Haines
Yahoo
haines@yahoo-inc.com

Mounia Lalmas
Yahoo
mounia@acm.org

ABSTRACT

The advertising industry is seeking to use the unique data provided by the increasing usage of mobile devices and mobile applications (apps) to improve targeting and the experience with apps. As a consequence, understanding user behaviours with apps has gained increased interests from both academia and industry. In this paper we study user app engagement patterns and disruptions of those patterns in a data set unique in its scale and coverage of user activity. First, we provide a detailed account of temporal user activity patterns with apps and compare these to previous studies on app usage behavior. Then, in the second part, and the main contribution of this work, we take advantage of the scale and coverage of our sample and show how app usage behavior is disrupted through major political, social, and sports events.

1. INTRODUCTION

The increasing usage of mobile devices and mobile applications (apps) has been a game changer in the technology industry. The emergence of online marketplaces and APIs has allowed developers, market intermediaries, and consumers to develop, disseminate, and use mobile apps. As a consequence, understanding user behaviours with apps has gained increased interests. For example, marketplace operators want to identify popular or problematic apps to then provide effective app recommender systems, whereas developers want to understand why their apps are liked or disliked by the users to improve the app design.

Studies of mobile app usage behavior vary in terms of sample size and the coverage of user activity. On one side of the spectrum are small sample studies that aim to provide a comprehensive description of the app usage behavior of a particular group of users. These studies are able to capture details of hourly, daily, or weekly behavior patterns by installing some kind of logging service on each device e.g. [1, 2, 5, 9]. On the other side, larger sample studies provide insights in particular aspects of app usage behavior, such as number of installations or aggregated network traffic statistics e.g. [4, 6, 7, 12]. However, the usage patterns observed in

these studies are limited to the aggregated server side information available through for example an app store or the network carrier.

Our study of app usage behavior differs from previous work in terms of sample size and aim. Our *sample* of usage data is the largest studied so far as well as more comprehensive in terms of user activity coverage compared to any previous large scale studies. Our *aim* is to focus on behaviours that can be characterised as disruptions to typical mobile app behaviour as these may provide potential marketing and personalization opportunities.

Our contributions are the following. First, we provide a detailed account of the user activity patterns observed in the largest app usage dataset so far and where applicable compare these to previous studies. Then, in the second part, and the main focus of this work, we take advantage of the scale and coverage of our sample and show how app usage behavior is disrupted through major political (Brexit), social (new year's day), and sport (Euro 2016) events.

2. RELATED WORK

Some studies of app usage aim to infer various user demographics based on their app usage e.g. [5]. Others e.g. [12] identify distinct types of users based on their app usage, also demonstrating a strong relationship between demographics and app usage. For instance, by representing users as vectors with dimensions such as category, time of day, workday versus weekend, they associate meaning to some of the groups, such as evening learners, or screen checkers. We use similar dimensions in our work.

A large body of work focuses on gaining an understanding of actual app usage. Here we report some of the findings, as relevant to our work. Yang et al. [11] show that the average number of apps one user visits is 7 over the week, and 5 categories of applications within a day. Social network is the dominant app, then search, then e-commerce. Falaki et al. [1] found that users interact with their smartphone anywhere between 10 to 200 times a day on average and that session length varies between 10-250 seconds. They were unable to show a relation between longer sessions and having fewer or more sessions. Yan et al. [10] found that most interactions with the phone are short with 80% of the apps being used for less than two minutes, indicating that short engagement is the norm. Their analysis also reveals that users overwhelmingly use only one application per session. Yang et al. [11] also found that app usage is often focused at a given time; for instance during a particular time of day, users usually visit apps from a dominant application category (e.g. news app category in the morning).

Various periodic patterns were also identified. For instance, Xu et al. [8] found that the diurnal patterns of different categories of apps can be remarkably different. For example, news apps are

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW'17 Companion, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4914-7/17/04. <http://dx.doi.org/10.1145/3041021.3051113>



much more frequently used in the early morning, sports apps are more frequently used in the evening, while other apps have diurnal patterns less visible and their usage is more flat during the day. Games apps also peak after standard work hours as we would expect, since that is probably the typical recreation time for most subscribers. Li et al. [4] also found that app usage changes during a day, and keeps growing from 6am and reaches the first peak around 11am, and declining slightly between 11am to 12pm, etc. They also find that about 32% of app usage are performed during 7pm to 11pm, reaching the maximum around 9pm. Such a distribution is quite consistent with human habit. After 9pm, app usage declines quite sharply, and reach the minimum around 5am. In our work we also show similar patterns.

Falaki et al. [1] found that heavy users tend to use their phone more consistently during the day whereas light users tend to have concentrated use during certain hours of the day. In terms of application usage, a large number of applications installed by the users does not mean that they use them equally. Users devote the bulk of their attention to a subset of applications of their choice. In addition, as shown in [10], location has an effect on app usage, for example, a tendency for game usage is at home whereas work-related app usage is at work.

All the above clearly suggest that app usage indeed follows *regular patterns*, in terms of which app, and when during the day or the week they are mostly used.

In a somewhat related line of work, studies on second-screen device use during television viewing have been carried out. For instance, Holz et al. [2] found that during television programs, users on average launched 0.06 apps per minute on their phones. This goes down to 0.04 during commercials. The median duration of app usage was 22 seconds on average for apps launched during programs and 19 seconds during commercials. They further saw that certain TV genres provoke substantially more app use than others, such as sports. Finally, contradicting previous published work, they found that none of the web pages or domains participants visited while watching TV related to the content of the programs being watched.

3. DATA CHARACTERISTICS

We obtain our sample of app usage data from Flurry,¹ a library that mobile developers integrate in their apps to measure app usage and allow in-app advertising.

An app that integrates Flurry logs a particular set of default app events triggered by user actions such as an *app start* event on opening an app or an *ad click* event when clicking on an advertisement. App developers may define app-specific events. For example, for a game app a level-up event could be created and would subsequently be logged. All events are stored on the server side and provide a rich source of information for the analysis of user behavior with apps.

Since apps differ in the events they log as well as in the user behavior that they elicit, e.g., news vs. e-mail app, we focus on user engagement metrics to characterize user behavior. That is we focus on discovering patterns in user engagement with mobile apps over time instead of specific app event patterns.

There are three types of engagement measures: (i) popularity (e.g. number of users, number of sessions); (ii) loyalty (e.g. number of active days per user); and activity (e.g. time spent, number of clicks) [3]. We use a popularity based engagement metric and measure the number of sessions a user has with an app based on the app start event. Each time a user opens an app, either a new session

¹<https://developer.yahoo.com/analytics/>

device feature	%session	%apps	%users	user feature	%session	%users
OS: android	78%	38%	67%	gender: female	54%	55%
OS: iOS	21%	59%	32%	gender: male	46%	45%
Make: Samsung	36%	32%	31%	age: 13-17	9%	7%
Make: Apple	21%	59%	32%	age: 18-24	18%	17%
Make: LG	6%	17%	4%	age: 25-34	19%	17%
Make: Sony	3%	16%	3%	age: 35-54	47%	46%
Make: Motorola	3%	14%	3%	age: 55+	6%	12%

Table 1: Coverage statistics of device (left) and demographics (right) information.

is recorded or a previous session is resumed. If the user leaves an app, but revisits within 10 seconds (without visiting other apps) the session resumes; otherwise the session ends and a new session is recorded upon opening the app.

We collected a sample of Flurry data from May 2016, consisting of events from 230K mobile apps and 600M daily unique users. Our sample covers users from 221 countries. 61% of the apps were accessed from the US and 34% from the UK. In the subsequent sections, we focus on these two countries.

In addition to events, meta information about the app, e.g., category, the user of the app, e.g., age, and the device the app is installed on, e.g., operating system, is recorded. Table 1 (left) shows the characteristics of the various mobile phones on which apps using Flurry are stored. 99% of the devices use either Android or iOS as operating system (OS); the remaining 1% consists of Windows and Blackberry. Further, 67% of the users use android while 32% uses iOS. This roughly aligns with figures reported in NetMarketShare (May, 2016) that reports a market share of 71% for Android and 23% for iOS.²

When registering with an app, users may be asked to provide information about their gender and age. These demographics are recorded in Flurry. In addition, for apps with no such information (e.g., when apps do not solicit such data), Flurry infers the demographics of a user on the basis of the other apps that same user has on his or her mobile device. Table 1 (right) shows the percentage of session starts, and daily users per age and gender category. Most users are between 35 and 54 year olds (46%), and there are more female users than male users (55% vs. 45%). The percentage of sessions is generally proportional to the number of users, which aligns with results reported in [3], i.e. more users result in more sessions. However, for users aged 55 and above, the percentage in sessions is half that of the number of users suggesting that these users launch apps less frequently.

Each app incorporating Flurry registers an app category. Our dataset contains 27 categories ranging from work related (e.g., productivity) to leisure (e.g., games) and other popular categories (e.g., news). The most popular app categories, in terms of number of daily sessions, daily visited apps, and daily unique users, as well as categories with the highest average session length, are listed in Table 2. There is a high number of game apps (33%), however these apps contribute to a smaller proportion of the sessions (12%) than social or utility apps. However, the session length for game apps is high (6 minutes) and about twice the session length of an average app. This suggests that users do not start games apps as often as for example social apps, but when they do a longer time is spent within the app. We make a similar observation for news and e-reader apps; when users open a reading app, they also spend a longer time within the app, e.g., 7 minutes average session length.

²<https://www.netmarketshare.com/operating-system-market-share.aspx?qprid=8&qpcustomd=1&qpssp=208&qptimeframe=M>

This finding is contrary to those by Falaki et al. [1] who found no relation between session length and session starts; however, their sample of 255 users showed large diversity which may have hidden this relation. Further, the age distribution is different from the one reported in other work [4, 12] likely due to the difference in user base samples (worldwide versus a specific country). Other figures are in accord with those reported in e.g. [4, 12] in particular in terms of gender and app popularities.

4. USER ENGAGEMENT PATTERNS

We now explore how temporal app engagement patterns vary for different app categories and demographics. The goal is to provide a purely observational description of user engagement patterns. This serves as a basis for our experimental analysis of disruptions in Section 5. We use the number of session starts as our proxy of user engagement and aggregate these for a particular time unit. This allows us to measure users’ engagement with apps at a particular time interval. We focus on users in one country, the US, and account for different time zones.




Figure 1: Normalized number of sessions during an average weekday, resp. weekend day in the US.

General daily engagement patterns. Figure 1 shows the normalized average number of sessions per hour for week days and the weekend as purple (solid) and yellow (dotted) respectively. Mobile activity decreases during the night and increases during the day. Users become more active on their mobile device during the day and are most active during the evening (around 9pm). This corresponds with findings in previous work by Li et al. [4]. Figure 1 shows that on weekdays engagement is higher in the morning than on weekends. In the weekend users are more active till later in the night than during the week and become active later in the morning. This may be explained as people going to bed later and sleeping longer during the weekend.

Daily engagement patterns by category. The general engagement pattern with apps shows a clear increasing trend during the day, peaking in the evening. Next we ask whether engagement patterns vary with apps belonging to different categories. We are particularly interested in identifying the app categories for which engagement patterns deviate from the average pattern for all apps. To identify these categories, we use the cumulative absolute difference

	%session	%apps	%users	avg. session length			
1. Utilities	23%	Games	33%	Games	25%	E-readers	7m
2. Social	16%	Entertainment	8%	Utilities	22%	Health & Fitness	6m
3. Games	12%	Lifestyle	6%	Social	17%	Games	6m
4. Productivity	11%	Productivity	6%	Productivity	13%	Entertainment	5m
5. Personalization	9%	Education	5%	Photography	7%	News	5m

Table 2: Coverage statistics for top 5 app categories.




Figure 2: Normalized number of sessions during an average week day in the US for different app categories.

(CAD) between the time series of a category and the average time series:

$$CAD(c) = \sum_{i \in [0,23]} |s_{a,i} - s_{c,i}| \quad (1)$$

where c is an app category, $s_{a,i}$ is the normalized average number of sessions at hour i , and $s_{c,i}$ is the normalized average number of sessions for category c at hour i .

The app categories that differ most from the average pattern are education, finance, sports and weather. Figure 2 shows the normalized average number of sessions for each category during weekdays. Educational and finance apps are more popular during the day (between 7am and 4pm) than the evening (after 5pm). The usage of weather apps peaks in the morning (between 5am and 9am), as users check the weather before starting their day or going out. Sport apps become popular in the evening (between 6pm and 10pm) as users engage in sport related activities (going to the gym, watching a game) from that time onwards (e.g. after work). This is inline with previous findings [8].

We also consider whether app engagement patterns differ per day of the week. Figure 3 shows the normalized average number of sessions for each day of the week. School and work related apps (education and finance) are most popular during the week, whereas leisure-related apps (sports) are most popular during the weekend.




Figure 3: Normalized number of sessions during an average week in the US for different app categories.

This section shows that app usage follows clear temporal patterns in terms of app engagement. These patterns are intuitive, and are in accord with patterns identified in other studies e.g.[1, 8, 11]. The results also serve as a sanity check for our dataset as to the presence of any bias.

5. IDENTIFYING DISRUPTIONS

Previous mobile app usage studies have generally focused on identifying user behavior patterns like those described in Section 4.

The challenge of identifying whether and how user behavior is disrupted has received less attention, primarily due to two reasons: (i) disruptions are often rare and unusual events, therefore to detect disruptions either a sample over a prolonged period of time or a sample focused around a specific time known to contain a disruption is required; (ii) when a disruption is identified in a particular sample it is hard to attribute it to a particular external factor (or vice versa), as they may coincide by chance. However, the larger the sample the less events are likely to have a disruptive impact.

As discussed in Section 3, the scale and coverage of our sample make it uniquely suited for the study of disruptions in app usage behavior. We take a confirmatory approach to our analysis and define the following expectations regarding disruptions of user behavior during major political, sports, and social events: (i) app engagement during half time and after a Euro 2016 Championship match is higher than on an average day; (ii) app engagement with financial and news apps is higher in the days directly following the UK June 26 Brexit vote than on an average day before or some time after the vote; (iii) app engagement with social apps is higher on new years day than on an average day.

We describe our methodology to identify patterns of disruption during a target event, then describe our three use cases and determine if and how major events disrupt mobile engagement patterns.

5.1 Methodology

To determine whether mobile engagement changes more during a particular event than expected, we require an estimate of typical engagement patterns. We define typical behavior by averaging the mobile engagement over several instances of a day that does not include the target event. These days may be occurring before, after or surrounding the target day depending on the type of event and what is reasonable given a target day. For example, if the target event occurs on a Saturday then we take a number of Saturdays before the event. We refer to these non-event days as the “reference days”. Reference days themselves may contain anomalies that cause the total number of sessions starts to be unusually high or low. Reasons for this may include outages, new app releases (push notification), or other major events causing a change in user behavior. Days identifies as containing such outliers are removed from the set of “reference days.” An outlier, in this case a date, is defined as one for which the number of start session events is either lower than $Q_1 - 1.5 \cdot IQR$ or larger than $Q_3 + 1.5 \cdot IQR$, where IQR is the interquartile range of the number of start session events, Q_1 is the first quartile and Q_3 the third quartile.

To decide if an event may disrupt mobile usage, we divide a period (e.g., a day) into different time segments (e.g., 15 minutes) and normalize the number of sessions for each of these time segments. The expected number of sessions per time segment t is then estimated by averaging the normalized number of sessions for the reference days, noted by avg_t . In addition, we compute the standard deviation std_t for the reference days. We define behavior as “significantly” disrupted if the normalized number of sessions during the target period is larger than $avg_t + 2 \cdot std_t$ or smaller than $avg_t - 2 \cdot std_t$.

5.2 Use Case: Euro 2016

In our first use case, we consider whether matches (games) of the EURO 2016 UEFA European Championship (Euro 2016) coincide with disruptions in mobile app engagement. If so this would indicate a potential relation between the event and the disruption. We expect mobile engagement to be lower during matches while people are watching the game and higher during half time as well as right after the match as people catch up on missed notifications or

time	countries
Sat. June 11, 2016 20:00 GMT	England 1 - 1 Russia
Thu. June 16, 2016 14:00 GMT	England 1 - 1 Wales
Mon. June 20, 2016 20:00 GMT	Slovakia 0 - 0 England
Sun. July 10, 2016 20:00 GMT	Portugal 1 - 0 France

Table 3: Considered games in Euro 2016.

share their thoughts on the match. We examine this intuition using the following hypotheses: (H1) app engagement during a EURO 2016 match is lower than app engagement during the same time on an average day; (H2) app engagement during halftime of a EURO 2016 match is higher than app engagement during the same time on an average day; and (H3) app engagement in the 15 minutes after a EURO 2016 match is higher than app engagement during the same time on an average day.

The Euro 2016 matches were widely watched online and on television. In the UK over 16M viewers (25% of UK population) watched Portugal beat France in the final on BBC One. The event drew about 59% of the audience available at transmission.³ We focus our analysis on the UK, the country in Europe with most daily sessions in our dataset, and changes during EURO 2016 games. Table 3 shows the games used in our analysis. We model the typical mobile engagement for a match played on Saturday as the average engagement of all Saturdays between November 2015 and June 2016. We repeat this process to model typical app engagement on the reference days that are counterpart to each of the match days, for instance, Thursdays and Mondays. Each event day has 30 reference days.

Figure 4a shows the normalized number of sessions per 15 minute slot, on Saturday June 11, 2016 (England – Russia). The game starts at 20:00 and ends at 21:45 with a half-time break between 20:45 and 21:00. The two parts of the game are indicated by gray vertical lines. The green bars indicate the expected number of sessions, i.e., the average sessions starts over the reference days, and the blue lines indicate two times the standard deviation. The dots indicate the normalized number of sessions on June 11, 2016. The number of sessions $s \in [av - std, av + std]$ are indicated in green, the number of sessions $s \in [av - 2 \cdot std, av - std]$ and $s \in [av + std, av + 2 \cdot std]$ in yellow, and the number of sessions $s < av - 2 \cdot std$ and $s > av + 2 \cdot std$ in red.

Regarding H1 we observe that app engagement during games is not lower than during the same time on an average day for any of the matches listed in Figure 4. Contrary to our expectation, during matches, people appear to be as engaged with their phones as they would be during a normal day.


Regarding H2 we observe that during half-time app engagement is significantly higher than during the same time on an average day for the England – Wales match (at 14:45 - 15:00 Fig. 4b) and the England – Slovakia match (at 20:45 - 21:00 Fig. 4c). During the other matches we find engagement is up more than one standard deviation; however, this change is not significant.

Regarding H3 we observe that during the 15 minutes after each match, app engagement is significantly higher than during the same time on an average day consistently for all matches.


These findings are consistent with our hypotheses that people are more engaged with their mobile phone during half-time and shortly after a match. However, they are contrary to the observations in a study of second screen use during television viewing [2], where

³<http://www.bbc.co.uk/sport/football/36755192>


users used more apps during the show than during commercials. The sample used in that study, however, was small (seven households) and not all family members were equally interested in the programs watched.




(a) England - Russia, Saturday June 11, 2016



(b) England - Wales, Thursday June 16, 2016.



(c) Slovakia - England, Monday June 20, 2016.



(d) Portugal - France, Sunday July 10, 2016.

Figure 4: Number of UK sessions during selected EURO 2016 games.

5.3 Use Case: Brexit

In this use case, we consider if a large political event disrupts the app behavior patterns of mobile users. In particular, we focus




Figure 5: Number of UK sessions during day of the Brexit referendum outcome June 24, 2016.

on the Brexit referendum outcome in the UK and formulate the following hypotheses: (H4) finance app engagement directly after the Brexit referendum is higher than app engagement on an average day before the referendum; and (H5) news app engagement directly after the Brexit referendum is higher than app engagement on an average day before the referendum.

The United Kingdom European Union membership referendum, also known as Brexit, took place on Thursday June 23, 2016 in the UK to gauge support for the country's continued membership of the European Union. The result was announced in the early morning of June 24, 2016. The referendum resulted in an overall vote to leave the EU, by 51.9% on a national turnout of 72%. This outcome resulted into instability in the financial markets as well as turmoil in the UK political landscape. It is expected to take several years to decide how and when the UK will leave the EU.

We examine whether the announcement of the outcome on Friday June 24, 2016 coincides with disruptions in the mobile behavior of UK users. As reference days we take all weekdays in June before June 24. Figure 5 shows the percentage of sessions for different app categories on June 24 compared to the average usage. Regarding H4 and H5 we observe that the largest increases in app engagement are with finance (114%) and news apps (43%) and that this increase in engagement is significantly higher than app engagement on an average day before the referendum. Further magazines apps also show a significant increase in engagement of 24%. The increase in the popularity of these app categories, resulted in a decrease of the popularity of other categories such as lifestyle, health & fitness and entertainment. These findings are consistent with our hypotheses that people become more engaged with news and finance shortly after a large political event.

In addition to the short-term impact of the Brexit referendum outcome on news and finance apps, we are also interested in the long-term impact on user behavior. In particular, we are interested in the longevity of the increase in engagement with the finance apps, which were most impacted after the day of the referendum. Therefore, we plot the usage of finance apps in percentage of sessions throughout June to August 2016 in Figure 6. We observe a pattern of stable usage during the week and a drop during the weekend as well as a peak on June 24th, the day the referendum result was announced. We further observe a long-term effect in the usage of finance apps during the four weeks after the referendum announcement compared to the weeks before. The usage comes back to a similar volume as before the announcement at the end of July.




Figure 6: Percentage of UK sessions of finance apps during June, July and August 2016.

5.4 Use Case: New Year's Day

New Year's Day, the first day of the new year, is observed in most Western countries on January 1. Common traditions include attending parties, making resolutions for the new year, watching fireworks displays and calling one's friends and family. In this use case, we examine whether New Year's Day coincides with disruptions in mobile engagement patterns. Here, we focus on users located in the United States, the country with the largest coverage in our dataset. The week days between December 15, 2015 and January 15, 2016, without January 1, are used as reference days.

Figure 7 shows the percentage of sessions for the 10 categories with the largest percentage change in app engagement on January 1 compared to an average day. Photography and social are the only two categories whose usage significantly increases on New Year's Day, where in addition, the percentage usage of the other categories decrease. Photography usage increases with 40%, and the usage of the social category increases with 6%. Users appear to take a lot of pictures with their mobile device on New Year's Day (e.g., at the party they attend, of friends and family, etc.) and perhaps share on social media. Social messaging platforms are also often used to contact friends and family. The categories with the largest percentage decrease are education (-63%), finance (-57%), weather (-37%) and health & fitness (-28%). Users appear uninterested in these apps during New Year's day.




Figure 7: Number of US sessions during New Year day on January 1, 2016.

In this paper we studied user behaviours in terms of app engagement patterns and disruptions of those patterns in a data set unique in its scale and coverage of user activity. We confirmed some of the temporal user behavior patterns reported in previous work and highlighted cases where our findings deviate from those previously reported. These results provide direction for new questions regarding behavior with mobile apps, for example, about usage during commercial breaks on television or the relation between session length and number of sessions.

Unlike previous work, we explored the feasibility of identifying disruptions in app engagement behavior. Through three use cases

we demonstrate that changes in app engagement coincide with major sport, political, and social events. A more careful experimental design, however, is required to identify the strength of these potential relationships between events and app engagement.

In future work we plan to validate our findings through application in advertising and in personalization scenarios. For example, marketing campaigns targeted at specific app users that are more or less likely to engage with a specific event. Or identifying users that are more or less likely to be disrupted by particular events.

Acknowledgements: This work was performed as part of Steven Van Canneyt internship at Yahoo, London.

References

- [1] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in smartphone usage. In *MobiSys '10*, pages 179–194, 2010.
- [2] C. Holz, F. Bentley, K. Church, and M. Patel. I'm just on my phone and they're watching tv: Quantifying mobile device use while watching television. In *TVX'15*, pages 93–102. ACM, 2015.
- [3] J. Lehmann, M. Lalmas, E. Yom-Tov, and G. Dupret. Models of user engagement. In *UMAP'12*, pages 164–175, 2012.
- [4] H. Li, X. Lu, X. Liu, T. Xie, K. Bian, F. X. Lin, Q. Mei, and F. Feng. Characterizing smartphone usage patterns from millions of android users. In *IMC'15*, pages 459–472, 2015.
- [5] E. Malmi and I. Weber. You are what apps you use: Demographic prediction based on user's apps. In *ICWSM*, 2016.
- [6] A. Patro, S. Rayanchu, M. Griepentrog, Y. Ma, and S. Banerjee. Capturing mobile experience in the wild: a tale of two apps. In *CoNEXT'13*, pages 199–210. ACM, 2013.
- [7] P. Welke, I. Andone, K. Blaszkiewicz, and A. Markowetz. Differentiating smartphone users by app usage. In *UbiComp'16*, pages 519–523, 2016.
- [8] Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying diverse usage behaviors of smartphone apps. In *IMC'11*, pages 329–344, 2011.
- [9] Y. Xu, M. Lin, H. Lu, G. Cardone, N. Lane, Z. Chen, A. Campbell, and T. Choudhury. Preference, context and communities: a multi-faceted approach to predicting smartphone app usage patterns. In *ISWC'13*, pages 69–76. ACM, 2013.
- [10] T. Yan, D. Chu, D. Ganesan, A. Kansal, and J. Liu. Fast app launching for mobile devices using predictive user context. In *MobiSys'12*, *MobiSys '12*, pages 113–126, 2012.
- [11] J. Yang, Y. Qiao, X. Zhang, H. He, Fang, and G. Cheng. Characterizing user behavior in mobile internet. *IEEE ToETC*, 3(1), 2015.
- [12] S. Zhao, J. Ramos, J. Tao, Z. Jiang, S. Li, Z. Wu, G. Pan, and A. K. Dey. Discovering different kinds of smartphone users through their application usage behaviors. In *UbiComp'16*, pages 498–509, 2016.