A Large-Scale, Long-Term Analysis of Mobile Device Usage Characteristics

DANIEL HINTZE, FHDW University of Applied Sciences Paderborn
PHILIPP HINTZE, University of Münster
RAINHARD D. FINDLING, University of Applied Sciences Upper Austria
RENÉ MAYRHOFER, Johannes Kepler University Linz

Today, mobile devices like smartphones and tablets have become an indispensable part of people's lives, posing many new questions e.g., in terms of interaction methods, but also security. In this paper, we conduct a large scale, long term analysis of mobile device usage characteristics like session length, interaction frequency, and daily usage in locked and unlocked state with respect to location context and diurnal pattern. Based on detailed logs from 29,279 mobile phones and tablets representing a total of 5,811 years of usage time, we identify and analyze 52.2 million usage sessions with some participants providing data for more than four years.

Our results show that context has a highly significant effect on both frequency and extent of mobile device usage, with mobile phones being used twice as much at home compared to in the office. Interestingly, devices are unlocked for only 46 % of the interactions. We found that with an average of 60 interactions per day, smartphones are used almost thrice as often as tablet devices (23), while usage sessions on tablets are three times longer, hence are used almost for an equal amount of time throughout the day. We conclude that usage session characteristics differ considerably between tablets and smartphones. These results inform future approaches to mobile interaction as well as security.

CCS Concepts: • **Human-centered computing → Empirical studies in HCI**; **Mobile devices**; **Tablet computers**; **Empirical studies in ubiquitous and mobile computing**;

General Terms: Human Factors, Security, Measurement

Additional Key Words and Phrases: Daily interactions, Device unlocking, Locked usage, Session length, Smartphone, Tablet, Usage session, User context

ACM Reference format:
Daniel Hintze, Philipp Hintze, Rainhard D. Findling, and René Mayrhofer. 2017. A Large-Scale, Long-Term Analysis of Mobile Device Usage Characteristics. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2, Article 13 (June 2017), 21 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

Preliminary versions of this work have been published in UbiComp 2014 [9] and MoMM 2014 [10], which are extended by using an updated version of the underlying dataset twice the size, a more detailed analysis, and specific consideration of differences in device locking.
Authors' addresses: D. Hintze, FHDW, Fürstenallee 5, 33102 Paderborn, Germany; email: daniel.hintze@fhdw.de; P. Hintze, University of Münster, Malmedyweg 15, 48149 Münster, Germany; email: philipp.hintze@uni-muenster.de; R. Findling, University of Applied Sciences Upper Austria, Softwarepark 11, 4232 Hagenberg, Austria; email: rainhard.findling@fh-hagenberg.at; R. Mayrhofer, Johannes Kepler University Linz, Altenbergerstr. 69, 4040 Linz, Austria; email: rene.mayrhofer@jku.at.
 2474-9567/2017/6-ART13 $15.00
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Personal mobile devices have become ubiquitous today and people typically spend several hours using smartphones and tablet computers each day. Studying this symbiotic relationship between humans and personal mobile devices by analyzing the characteristics of user interactions with their devices can benefit many research areas. Examples are mobile data traffic prediction [24], indoor air quality monitoring [19], cognitive bias modification [23], compulsive behavior and technostress [17], smartphone addiction [15, 16, 18], healtcare [6], education [14], and user authentication [8, 34].

Consequently, smartphones – being the most popular mobile device form factor today – have recently been the subject of handset-based studies analyzing characteristics of usage and interaction [4, 5, 9, 10, 21, 25, 27]. However, little is known about how users interact with tablet devices, which are becoming a mainstream phenomenon, replacing traditional notebooks and desktops PCs in many areas. Smartphones and tablets offer comparable technical capabilities like connectivity, computational power, operating systems and application ecosystem. The two form factors differ predominantly in screen size. As device size has an effect on both application and mobility, understanding how tablets are used in comparison to smartphones is worthwhile. Based on the Device Analyzer dataset [30, 31], the largest mobile device usage dataset publicly available today, we therefore analyze mobile device usage characteristics such as session length, interaction frequency and daily usage with respect to three dimensions:

(1) As the majority of interactions with mobile devices do not include unlocking the device [9, 10], we distinguish between locked and unlocked usage.
(2) Since location context (e.g., being at home or at work) is suspected to have a noticeable effect on mobile device usage [25], we consider contexts classified as *home*, *office*, *other meaningful place*, and *elsewhere*.
(3) With little previous knowledge about the impact of form factor on device usage, this work is to our best knowledge the first to give a detailed comparison of usage characteristics for both smartphones and tablets.

Our objectives are two-fold: on the one hand we aim to give a high level overview of mobile device usage characteristics. On the other hand we want to provide extensive multi-layered statistical information on device usage based on the dimensions stated above. Considering three dimensions of mobile device usage, we seek to answer our main research question: *How do context, form factor, and lock status effect mobile device usage session characteristics?*

The paper is organized as follows: First, previous mobile device usage studies and their results are discussed in section 2. In section 3 we outline the underlying dataset and how usage sessions are derived, the algorithms applied to detect locations based on Wi-Fi scan results and GSM cell-IDs, and how contextual meaning is assigned to discovered locations. We introduce and discuss our findings in section 4 and explicitly describe current limitations in section 5. The final section 6 concludes the paper.

2 RELATED WORK

In recent years, a number of studies have examined different aspects of mobile device usage. Verkasalo [29] analyzed contextual patterns in mobile device usage based on usage logs from 324 smartphone users, finding device usage to be noticeably diverse in *office* and *home* context. Falaki et al. [4] examined user interaction on 255 Android and Windows Mobile smartphones and reported "immense diversity" in smartphone usage with the average number of interactions varying from 10 to 200. Oliver [21] conducted a large-scale but short-term (17 days on average) smartphone usage study on 17,300 BlackBerry devices, analyzing interaction time, interaction sessions and diurnal patterns. Böhmer et al. [1] captured application usage logs from 4,100 Android devices, observing that at night time the most popular applications are Facebook, Kindle, and Angry Birds. Soikkeli [25] studied the relation between mobile device usage and end user context based on usage logs from 140 smartphones.

The authors found usage sessions to be longer in *home* context while more frequent in *office* context. Most of the previous studies focused exclusively on smartphone usage. An exception comes from Müller et al. [20], who conducted a multi-method based exploration of tablet usage ($n = 33$), finding tablets to be mostly used at home and often while doing secondary activities such as watching TV, eating or cooking. Based on an earlier version of the Device Analyzer dataset used in our work, Wagner et al. [31] observed that a noticeably number of interactions occur without unlocking the device. The first work differentiating between locked and unlocked mobile device usage was conducted by Truong et al. [27], who conducted a small ($n = 10$) user study to analyze how often users unlock their devices. Finley and Soikkeli [5] examined multidevice usage (smartphone and tablet), observing that about 35% of multidevice sessions are dominated by a single device with only sparse usage of the second device. van Berkel et al. [28] provided a systematic model of smartphone usage, particular to analyze how researchers should handle brief gaps in interactions based on a field study with 17 participants. Harbach et al. [7] conducted a month-long field study with a panel of 134 smart phone users, focussing on the performance of different lock screen implementations, reporting that PIN users need more than twice as long before beginning the unlock process compared to users who use a pattern-based lockscreen.

Our study differs from previous work significantly in terms of duration and scale. With a mean of 144 days for phones and 230 for tablets, the sample period for devices in our analysis is higher than in any of the previous studies we are aware of. The total device usage time analyzed is 4,313 years, more than three times the extent of the time covered in [1], the largest handset-based mobile device usage study to our best knowledge (see table 3 for a comprehensive comparison).

3 METHODOLOGY

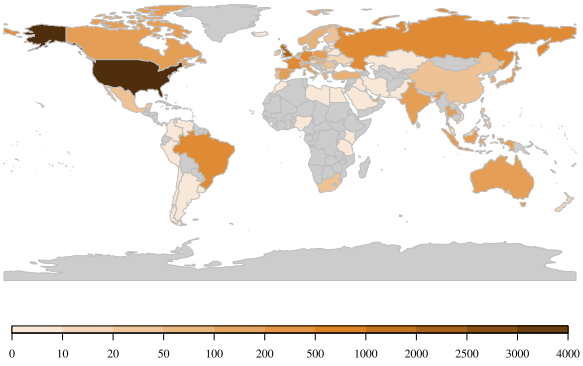

Fig. 1. Geographic distribution of devices within the dataset

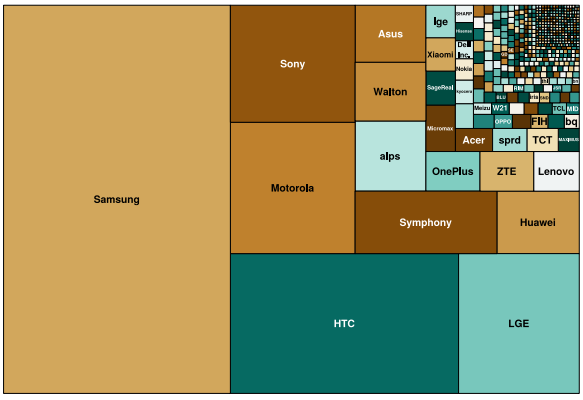

Fig. 2. Distribution of devices by manufacturer

3.1 Dataset

The analysis in this paper is based on the largest and most detailed dataset on Android device usage publicly available today, the result from the still ongoing Device Analyzer project [30, 31] by the University of Cambridge Computer Laboratory.[1] It consists of more than 225 billion records of Android mobile device usage, collected from 29,279 devices around the world. It captures 263 different features,[2] ranging from raw sensor data to application

[1]The University of Cambridge Computer Laboratory and Data Funder do not bear any responsibility for our analysis or interpretation of the Device Analyzer Dataset or data thereof.
[2]http://deviceanalyzer.cl.cam.ac.uk/keyValuePairs.htm

usage, recorded either periodically or event based by a stand-alone application available via Google Play Store. The dataset consists of 18 TB of log files, accessed using the Picky dataset sharing system [12].[3] Many devices within the dataset contribute data for an extended period of time, with 7,484 devices participating for more than one month, 535 devices providing data for more than one year and some even more than 4.5 years usage data. The dataset includes at least 1,277 different device types from 468 manufacturers (see fig. 2) and users from 175 different countries (see fig. 1). Since the Device Analyzer project emphasizes user privacy, no biographical or demographical features are available in the dataset.

To achieve the best data quality possible, we revised the dataset rigorously. Records created from older version of the client application which did not include all features required were disregarded. Only days captured entirely are used. Days during which usage was recorded only partially, e.g. due to crashes, application installation or deinstallation, or explicit pausing of the data collection were discarded. Because of this, out of 2.1 million days of device usage captured in the dataset roughly 1 million days (47.4 %) were disregarded. Days during which devices were powered on only for some hours were considered, as this might be part of the regular usage pattern. For day-based statistics we did, however, only regard days with at least one valid usage session.

Devices not using any keyguard were omitted, since they do not allow distinguishing between locked and unlocked state. We also removed devices configured to keep the display turned on while charging since this would distort the display state-based usage analysis. Finally, we only analyzed devices providing valid data for at least seven days. In total, these constraints led to the exclusion of 17,253 (58.9 %) out of 29,279 devices present in the dataset with a total of 1.3 million associated usage sessions.

In the last stage of the data filtering process, we excluded 1,493 devices for which we could not find a *home* context (see section 3.4), disregarding 2.7 million associated usage sessions.

The revised dataset used in this work contained 10,533 devices (9,861 phones and 672 tablets) with a total of 52.2 million usage sessions.

3.2 Usage Session Extraction

We consider mobile device usage sessions to be consecutive periods of time during which a user interacts directly with the device. Since mobile devices provide convenient access to their owners digital lives, they are typically protected against unauthorized access by some form of keyguard: for instance PIN, password, graphical pattern, face unlock, fingerprint, or swipe-to-unlock. While most of the device interactions require unlocking the device first, there are a number of restricted interactions possible without unlocking the device. The most common locked interactions are checking time, battery health, network connectivity, notifications, incoming calls, or taking pictures. Unlike most previous mobile device usage studies, we therefore distinguish between *locked usage sessions* and *unlocked usage sessions*.

Two different approaches to derive usage sessions from handset-based device monitoring logs have been used in previous studies. Since most device interactions involve the usage of an application, some authors [1, 5, 25] define usage sessions to be time intervals in which certain applications are running in the foreground of the devices. However, this approach is not suitable to study locked usage, as there is not necessarily an application active in the foreground during locked interaction. Mobile device interaction almost entirely relies on touchscreen interaction, either to display information or to capture user input. Because energy consumption is an inherent concern with battery powered mobile devices, displays – which are energy-intensive – are usually switched off as soon as possible after usage. This is done either manually or automatically after a short idle timeout. Hence, the more frequently used approach to derive usage sessions from device logs is to define usage sessions as time periods in which the device's screen is switched on (screen power based models) [4, 21, 22, 27].

[3] In this work we used a dataset snapshot generated on 16 May 2016 (Picky reference: `device_analyzer_full_20160516`).

Although naïve screen power based usage session extraction comes fairly close to actual device interaction, some pitfalls exist which – in our experience – can distort the results noticeably if not considered carefully. Consider e.g., incoming phone calls, which activate the screen to display the caller's number and to allow the user to answer the call. If the call goes unanswered, a naïve screen power based approach would falsely consider this a session of user interaction. Or consider phones with touchscreens that utilize a proximity sensor to switch off the screen when the device is held closely to a user's head, e.g., during a call, in order to prevent accidental touch events caused by the user's ear. As users tend to slightly shift the phone's position during calls, this would result in naïve screen power based models mistakenly recognizing multiple short usage sessions instead of one consecutive session. We observed that overall 12.7% of the changes in screen power state on smartphones are actually related to calls and hence do not constitute the boundaries of genuine user interaction sessions. These findings are based on a state machine based usage session extraction approach capable of avoiding mentioned pitfalls – which we consequently incorporate in our approach (see fig. 3).

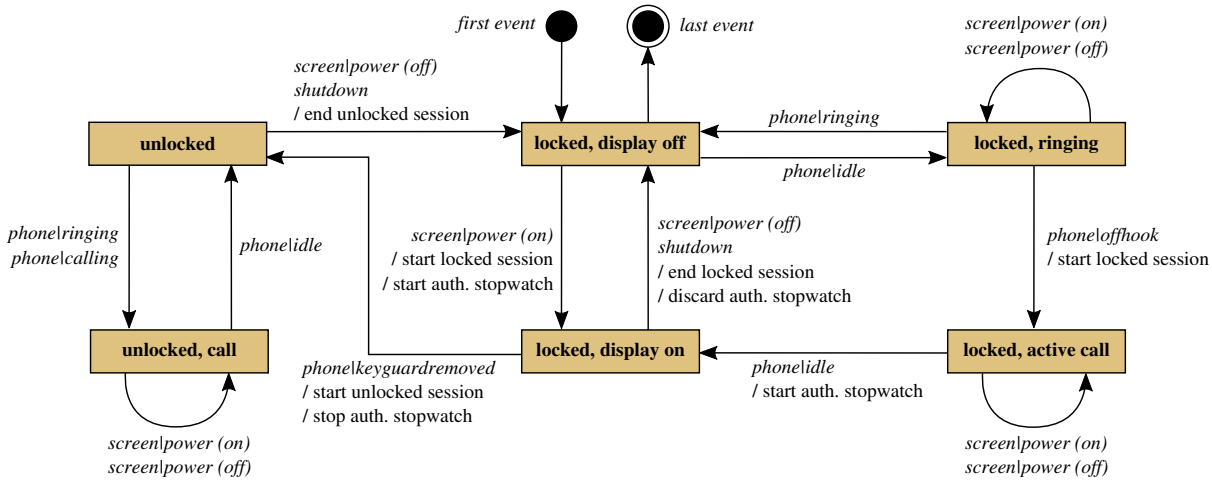

Fig. 3. State machine for session detection

3.3 Device Form Factor

We assume device form factors to have a considerable impact on device usage. We therefore analyze usage sessions characteristics with respect to device form factors – namely smartphone and tablet devices. One previously used approach to distinguish form factors in device logs is based on the device's ability to place or answer phone calls [9]. However, some tablet devices are capable of performing GSM voice call (e.g., the Galaxy Tab 10.1). Hence we chose the screen diagonal as a discriminator for form factors. Devices featuring a screen size of 7″ or higher are considered to be tablets while devices with smaller screens are regarded as smartphones. We calculate the screen diagonal from screen resolution and pixel density stated in the dataset.

3.4 User Context Classification

People use their mobile devices in different ways, depending on their current situation. For instance, in an office situation people might be more likely to use their smartphones to make phone calls or check for next meetings, while at home devices might be used more to browse the Internet or watch movies. Research by [25] reflects these

different usage patterns by observing that usage sessions are 37% longer in home context over office context, but happen 56% more often in office context over home context.

Deriving context from aggregated information is often difficult. Nevertheless, information on time and location can be combined in order to derive contextual place information. Based on previous research by [13] and [26] we distinguish four different place-related user contexts: *home*, *office*, *other meaningful*, and *elsewhere*.

While *home* and *office* are self-explanatory, *other meaningful* refers to places that do not have the characteristics of *home* and *office*, but still a significant amount of time is spent there. A frequently visited gym, for instance, would be considered an *other meaningful* place. Any place that is not classified as one of these three contexts is assigned the *elsewhere* context. This includes, but is not limited to, less frequent visited places like restaurants as well as transitions between other contexts.

Unlike other studies [13, 25, 26], we do not assign an *abroad* context for places outside users' home country. The reason being that [25] found that on average users spend only 2% of their time abroad, making this context negligible for the analysis of average usage patterns.

Alongside extracting locked and unlocked sessions, we derive the context these sessions occurred in, based on time and location information. While the Device Analyzer dataset provides timestamps, obtaining location information requires some effort. The dataset does not contain GPS information, which would be of little use in indoor or urban environments anyway. The Device Analyzer application records coarse locations of devices as returned by the network provider. Since recording such information raises privacy concerns, participants were requested to opt-in for sharing their location for research purposes – which only 1.12% of the users chose to do, precluding further analysis due to sample size.

Hence we derive location information from two other sources of information which can be related to device locations indirectly: GSM cell IDs and Wi-Fi scan results. GSM cell IDs were anonymized by hashing in the dataset to protect participants' privacy. Further, while the option to opt-out from recording anonymized GSM cell IDs existed too, only 2.41% chose to do so – leaving records for 97.59% of participating devices. Wi-Fi scan results, including SSID and MAC address of Wi-Fi access points within range are anonymized as well and are available for all capable devices in the dataset. An algorithm to extract location context information from handset-based GSM cell ID data has been proposed by [13], extended to utilize Wi-Fi scan data by [26] and applied to a study of smartphone usage in [25]. For our research, we implemented the extended algorithm while applying some simplifications for the sake of computation time ([25, 26] applied the algorithm on a dataset of 140 devices while the dataset we use contains 29,279 devices). The algorithm consists of two parts: first, meaningful locations are identified, which requires different approaches for cell ID data and Wi-Fi scan results. Subsequently, contexts such as *home* or *office* are assigned to the identified locations based on time information.

3.4.1 Deriving Places from Cell Data. A mobile phone is almost always connected to a cell tower, uniquely identified by cell identifier (CID) and location area code (LAC). As these attributes are anonymized in the dataset used in this work, we cannot relate them to geographic coordinates by using a database like OpenCellID[4]. However, since a cell tower has a fixed position and a limited range, it could be considered to be one place in terms of user context detection. As cell tower placement aims to minimize areas without network coverage and enhance connectivity robustness, adjacent cells usually overlap each other. Devices may dynamically switch between cells if another one is considered "better" than the current cell. As a result, it is not unlikely for even a stationary mobile phone to be connected to several different cells over the course of time [33]. Moving the device, for instance in an office building, possibly even increases the number of different cells a device is connected to while still being in the same abstract place (e.g., *office* context). In order to obtain places from cell data, adjacent cells therefore need to be clustered. For our implementation, we apply a clustering algorithm based on *minimum circular subsequences* proposed by [33]. Given a sequence of cell IDs a device has been connected to, [33] defines

[4]http://opencellid.org

a *circular subsequence* as a subsequence starting and ending with the same cell ID and containing at least two different cell IDs with the *cardinality* being the number of different cell IDs it contains. A *minimum circular subsequence* is a circular subsequence that does not contain other circular subsequences and thus indicates that a device has "returned" to where it was in the beginning. Cells that appear in a minimum circular subsequence of low cardinality are assumed to be co-located and therefore assigned to the same cluster. To avoid the problem of "over-clustering" large areas in situations like stop-and-go traffic on a freeway, cells are clustered around "qualified" cells that appeared at least Q times for at least one day. For our work, we choose $Q = 10$ and a minimum circular subsequence cardinality threshold $S = 2$, as suggested by [33]. Further details on deriving places from cell data are found in [13, 25, 33].

3.4.2 Deriving Places from Wi-Fi Scan Results. Wi-Fi-enabled mobile devices periodically scan for Wi-Fi access points within range. The result contains a list of access points, each described by its MAC address, SSID, RSSI, and frequency. The interval between individual scans ranges from a few seconds to several minutes, depending on factors like OS build, hardware, device state, and connectivity state. The dataset used in this work features an average scan frequency of 129 scans per day.

Since Wi-Fi access points are typically stationary, Wi-Fi scan results are frequently used for location-based services such as indoor positioning and navigation systems. A popular approach is to construct a unique Wi-Fi "fingerprint" of a certain location based on observed unique access point identifiers and corresponding signal strengths and an extensive body of literature exists on various fingerprinting techniques. While previous studies used a fingerprinting-based approach to derive meaningful places from Wi-Fi data [25, 26], we choose a less complex method. Taking the available history of scan results for a single device as input, the steps outlined in alg. 1 are applied to derive contextual places, each identified by a cluster of adjacent access points.

ALGORITHM 1: Wi-Fi Acess Point Cluster Algorithm

$A \leftarrow$ sequence of all known access points
$sort(A) \leftarrow$ sort descending by the number of occurrences.
while *A is not empty* **do**
  $R \leftarrow pop(A)$
  $C \leftarrow cluster(R)$ The first access point from A constitutes the root R of a new cluster C
  for *each access_point in scans_containing_R,* **do**
    $C \leftarrow C$ + *access_point*
    $A \leftarrow A - C$ Remove from A each access point contained in C
  end
end

While this approach is less sophisticated and presumably less accurate than a fingerprinting-based approach, it is also less complex and computationally intensive — an important factor for processing 18 TB of raw data on commodity hardware. Assuming a Wi-Fi access point has a maximum indoor range of 50 meters, a cluster spans at most a circular area with a diameter of 150 meters (imagine a cluster containing three access points with the root R located in the middle and the other two access points opposed to each other as far away as possible while still maintaining an overlap with R). As we are trying to identify places such as home and office (and keeping in mind that in contrast, GSM cells can have a range of several kilometers), we argue that the granularity of our approach is sufficient for the study at hand, allowing us to avoid a more computationally expensive fingerprinting-based approach.

3.4.3 Context Detection. Time information is one of the most important aspects available to detect user context [2]. Making some basic assumptions about standard users' diurnal patterns allows us to make a fair guess about *home* and *office* contexts: In order to put a contextual meaning to the places derived from cell and

Wi-Fi scan data, we assume that under normal circumstances a standard user does not sleep in the office, is at home during night hours (between 00:00 and 06:00), works between 10:00 and 16:00 on workdays, and does not regularly go to work on weekends.

While these assumptions are obviously fuzzy and oversimplified considering e.g., night shifts, home workers, holidays, unemployment, or traveling salesmen, previous research shows that results are still fairly accurate. Based on similar assumptions, [13] was able to detect *home* contexts with an accuracy of 66% and *office* contexts with an accuracy of 74% (n = 578) while [29] reported classifying 70% of contexts correctly (n = 87), both solely using places derived from cell information.

To detect *home* and *office* context we apply an algorithm based on [26] to both cell-based and Wi-Fi-based places. At first, places that have been visited more often than the average number of visits across all derived places are considered to be *meaningful* places. Places not classified as *meaningful* places are assigned the *elsewhere* context. Further, meaningful places are considered to be *office* context if both

$$\frac{\text{Visits during weekends}}{\text{Total visits}} < 0.2 \tag{1}$$

$$\frac{\text{Visits during weekday working hours}}{\text{Visits during weekdays}} > 0.5 \tag{2}$$

Home context is assigned to meaningful non-office places if both

$$\frac{\text{Visits on weekday night hours}}{\text{Visits during weekdays}} > 0.25 \tag{3}$$

$$\frac{\text{Visits on weekdays during non-working hours}}{\text{Visits during weekdays}} > 0.7 \tag{4}$$

Other meaningful is assigned to all meaningful places neither considered *home* or *office*.

Cell-based and Wi-Fi-based context detection is applied to classify the context of a usage session, depending on which information is available. If for one place contexts derived cell-based and Wi-Fi-based differ, we choose the most specific context in the following order: *home*, *office*, *other meaningful*, *elsewhere*.

3.4.4 Context Classification Performance. The Device Analyzer dataset does not contain any ground truth regarding location context, so we can not directly assess the accuracy of our context detection algorithm. We therefore applied the algorithm to a different, much smaller dataset that contained location labels. The dataset used for context classification performance assessment is the AlgoSnap Crowdsignals[5] pilot dataset, the result of a crowdfunded, handset-based data collection from August to November 2016 with 31 participants, 20 males and 11 females, of varying age, education and ethnicity. Of the 31 participants, 23 reported to be employed while 8 stated to be not employed, not able to work or retired. The dataset captures a variety of different features, including Wi-Fi scan results and cell connections. More importantly, participants where asked different labeling questions at random when unlocking their phone, one of which was asking for their current location. Allowed answers included, among others, Home, Work, Bank, Hotel, Church, and Restaurant, with the most recent selection being preselected. Participants could always choose to dismiss the question or disable them permanently but where paid $ 0.05 per response.

We applied the user context classification described above to the Crowdsignals dataset to measure the accuracy of the context prediction. We were not able to detect a *home* location for three of the participants, so they where excluded from further analysis (note that we also exclude devices for which no *home* context could be established from the device usage study). With the average number of context ground truth from the lockscreen survey being 757 labels per user, we also excluded five users who provided less than a quarter of the mean number of

[5]https://crowdsignals.io

ground truth labels. Finally, we removed three more users who seemed to have provided a significant amount of implausible labels. One user, for example, reported being in a restaurant every single time over the course of several weeks - most likely submitting the answer to earn money but not actually adjusting the selected location.

Table 1. Context detection confusion matrix

	Reference		
Prediction	Home	Office	Other
Home	3,777	704	1,616
Office	261	939	649
Other	1,494	854	1,735

Table 2. Classifier statistics by class

	Home	**Office**	**Other**
Sensitivity	0.6828	0.3761	0.4338
Specificity	0.6429	0.9045	0.7076
Pos Pred Value	0.6195	0.5078	0.4249
Neg Pred Value	0.7041	0.8470	0.7150
Balanced Accuracy	0.6628	0.6403	0.5707

For the remaining 21 users we evaluated the performance of our context classification algorithm against the ground truth labels from the Crowdsignals dataset. Aberrant from the algorithm stated above, we did not distinguish between *other meaningful* and *elsewhere* contexts because the ground truth labels did not allow to reliably distinguish them. The results in table 2 show that with 62% the performance of *home* context detection is close to [13], who reported a positive predictive value (PPV) of 66%. *Office* detection, however, performs less well with a PPV of only 51% whereas [13] reported 74%. As table 1 shows, class distribution in the ground truth is notably skewed, so the balanced accuracy offers a more meaningful metric. The classification achieves a balanced accuracy of 66% for *home* detection and 64% for *office* detection. When interpreting these results, however, one should keep in mind that the labels used as ground truth itself contain a certain degree of human error, thus limiting the validity of this performance analysis.

4 RESULTS AND DISCUSSION

In this section we present and discuss our results. Studying locked and unlocked usage sessions for certain characteristics constitute the core results of this work. Examined characteristics include: average device usage time per day, average usage session duration and average amount of usage sessions per day. For each locked, unlocked, and overall usage sessions we compute mean and median number of daily interactions as well as mean and median daily usage time in regard to context and form factor. For each device, this is done by calculating the mean and median for each feature over all observed days. The mean and median locked, unlocked, and overall session durations are calculated across the entire observation period for each device, again in relation to context and form factor. We then calculate the grand mean (mean of the means of all devices) and the grand median (median of the medians of all devices). Table 3 summarizes our results and compares them to findings of previous mobile device usage studies.

4.1 Context Detection

Comparing GSM and Wi-Fi-based location detection, as expected we found Wi-Fi-based location detection to yield better results in most situations. Quality of results was measured by the amount of distinctly detected home and office contexts. We assume two reasons for the higher quality of Wi-Fi-based location detection: First, Wi-Fi signals have a smaller signal range compared to GSM signals, hence allowing a more precise detection of locations. Secondly, parameterizing the clustering of cell IDs means to balance under- and over-clustering, in which either multiple clusters exist for one abstract location or multiple locations are falsely grouped together. Moreover, cell ID information are not available for around half the analyzed tablet devices. However, Wi-Fi-based location detection as well does not always yield results, for instance at work places without any Wi-Fi access points in range.

Hence, combining both location sources improved the overall result in every situation. In particular, *home* context could be detected for 88 % of the phones and 83 % of the tablet devices, as outlined in fig. 4. For 80 % of the phone-type devices *office* context was detected while only for 53 % of the tablet devices an *office* context was found. This is to be expected, considering that tablet devices are less handy and thus less often brought to work, compared to smartphones. To not distort results, we excluded devices for which no *home* context could be detected from consecutive usage session analysis.

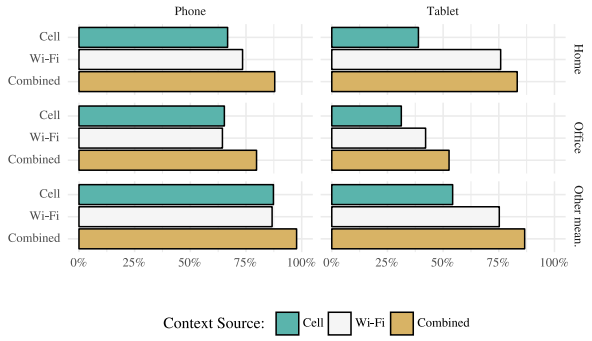

Fig. 4. Context detection results

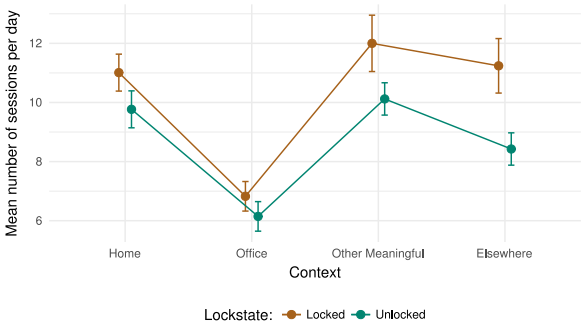

Fig. 5. Repeated measures ANOVA descriptives of daily session count on phones by context

4.2 Number of Daily Interactions

When looking at the number of daily interactions, it is noticeable that the majority of interactions does not include unlocking the device. Overall, people used their phones on average 60 times per day but only unlocked them for half (46 %) of the interactions. Tablet devices are used less than half as often, namely 23 times per day on average with a similar unlocked usage share of 38 %. Since locked usage only allows for a limited set of actions, mainly checking information, the high proportion of locked sessions can be explained by *checking habits* as described by Oulasvirta et al. [22]. We note that the average number of daily device interactions varies considerably across users, as figs. 6 and 7 illustrate.

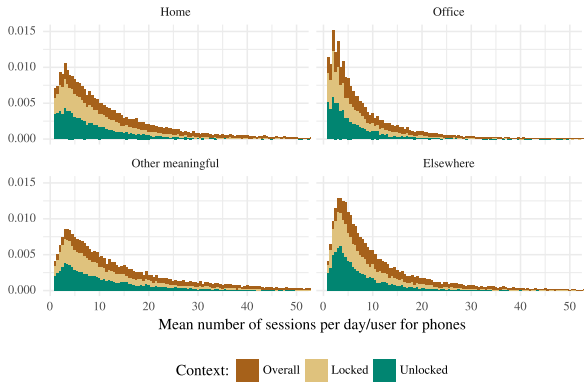

Fig. 6. Mean number of sessions per day/user on phones

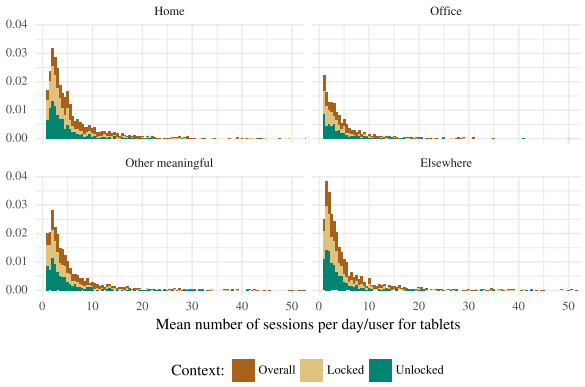

Fig. 7. Mean number of sessions per day/user on tablets

In the distribution of interactions across the different contexts (see figs. 8 and 9), we observe that on average, more device interactions occur in *other meaningful* places than at home while *office* has the fewest interactions, indicating that people use their devices less frequently in work situations compared to leisure activities.

Our results with respect to *office* usage are well in line with findings by Soikkeli [25], who reported that 12% of smartphone usage sessions occur in *office* context and 29% *elsewhere*. Our results indicate that the share of smartphone sessions in *office* context is 19 % while 22 % occur *elsewhere*. However, Soikkeli [25] found that 47% of the sessions take place in a *home* situation while *other meaningful* places only account for 9% of the sessions. We found, though, the share of phone sessions in *home* context to be 27 % while *other meaningful* places accumulate 32 % of the usage sessions. This effect might be introduced by different user panels: the Device Analyzer dataset we use contains users from 175 different countries and is not limited to specific professions, age groups, or life styles, while the panel used in Soikkeli [25] consists mainly of Finnish male university students.

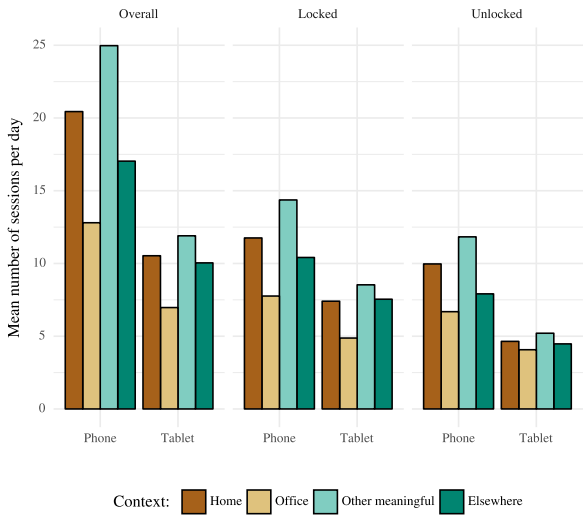

Fig. 8. Grand mean of number of sessions per day by context

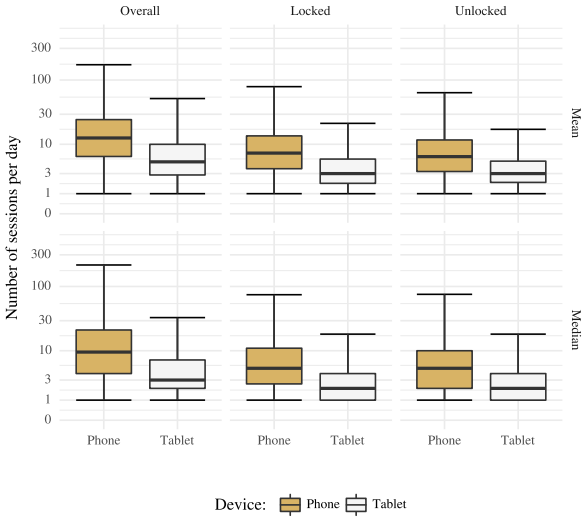

Fig. 9. Distribution of sessions per day across devices

To verify whether these trends indicate significantly different device usage with respect to context and lock state rather than noise in the data used, we calculated a repeated measurement ANOVA with the factors context (*home*, *office*, *other meaningful place*, *elsewhere*) and lock state (*locked* or *unlocked*) and observed a highly significant main effect of context, $F(2.86, 2650.03) = 57.749, p < 0.001$. Mauchly's test of sphericity indicated a violation of the assumption of sphericity ($p < 0.05$), so a Greenhouse-Geisser correction was used. Furthermore, there is a highly significant main effect of lock state, $F(1, 926) = 28.31, p < 0.001$. The interaction between context and lock state is also highly significant, $F(2.53, 2342.27) = 8.85, p < 0.001$. Again, Mauchly's test of sphericity indicated a violation of the assumption of sphericity ($p < 0.05$), so a Greenhouse-Geisser correction was used. Tuckey post hoc tests revealed significant differences in the average number of sessions at home compared to in the *office* ($p < 0.001$). No significant differences were found between the average number of sessions at *home* and in *other meaningful* places ($p = 0.47$) and between *home* and *elsewhere* ($p = 0.85$). Furthermore, significant differences in the average number of sessions in the *office* compared to *other meaningful place* ($p < 0.001$) and between *office* and *elsewhere* ($p < 0.001$). Likewise, significant differences in the average number of sessions in *other meaningful places* and *elsewhere* were found ($p = 0.008$). Figure 5 outlines the corresponding ANOVA descriptives.

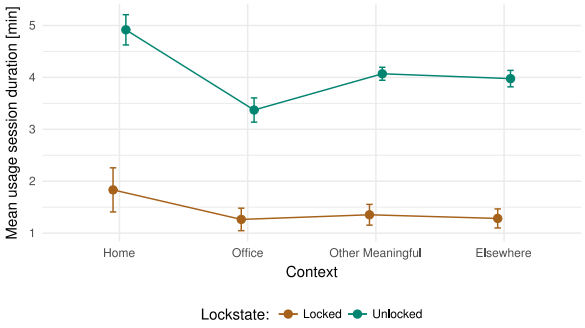

Fig. 10. Repeated measures ANOVA descriptives of mean usage session duration on phones by context

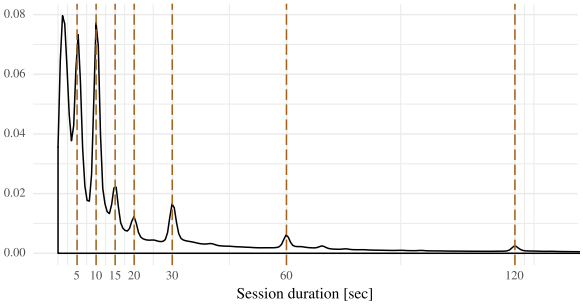

Fig. 11. Kernel density estimate for locked session duration distribution.

4.3 Session Duration

Regarding session duration, we found that in general, usage sessions on tablet devices last more than twice as long as phone usage sessions. Locked sessions on average last 107 seconds on phones (median 57 seconds) while spanning 271 seconds on tablet devices (median 84 seconds). Locked sessions being longer for tablet devices compared to smartphones are presumably caused by the fact that tablets are configured with an average display timeout of 6.6 minutes while smartphones feature a mean display timeout of only 2.8 minutes. As locked usage sessions are usually short, they are more prone to distortion caused by display timeouts counted towards usage time in cases in which the user does not manually switch off the device's screen, which technically marks the end of the usage session. Figure 11 illustrates the degree of distortion, outlining common display timeout intervals in the kernel density estimate for the distribution of locked session duration.

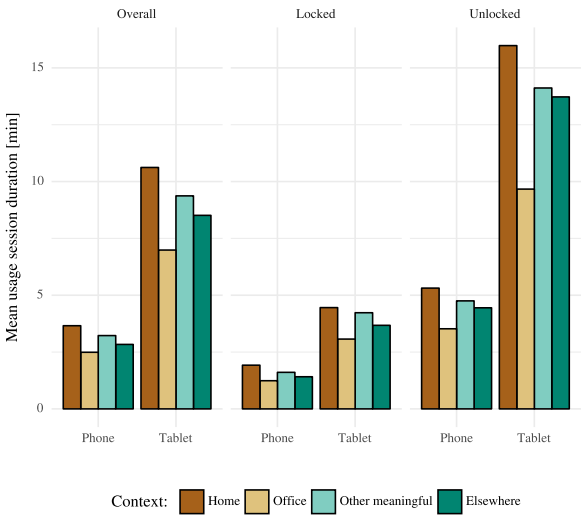

Fig. 12. Grand mean of session duration by context

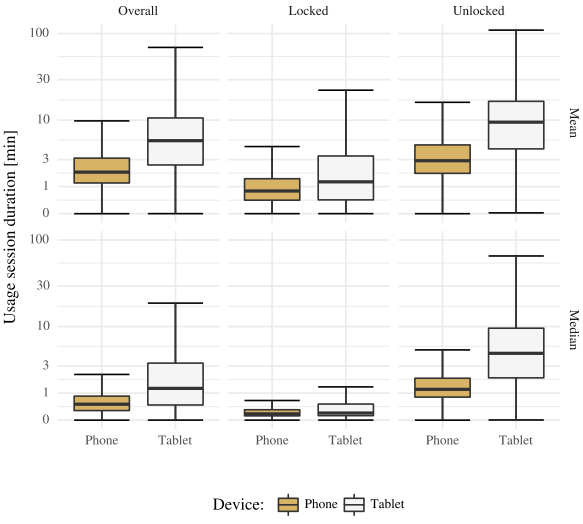

Fig. 13. Distribution of session duration across devices

Average unlocked sessions span 307 seconds on phones (median 73 seconds) while lasting for 963 seconds on tablets (median 297 seconds). Interestingly, context seems to have a noticeable effect on session duration (see

fig. 12): In *home* context, sessions on both tablet and phone devices are considerably longer than in other contexts while sessions in *office* context are usually the shortest. On tablet devices, for instance, unlocked sessions in home context have an average duration of 11.4 minutes while in office context, unlocked sessions would only last 6.7 minutes.

Again, session duration is highly diverse, both across sessions and across users by more than an order of magnitude. For example, the median across the mean unlocked session lengths of tablet devices is 683 seconds, compared to a mean of 963 seconds, which is biased by a mean session length of up to 387 minutes on some devices. Figure 13 therefore again depicts the distribution of both mean and median of the session duration per device for locked, unlocked and overall usage.

To verify the statistical significance of the observed trends a repeated measurement ANOVA with the factors context (*home*, *office*, *other meaningful place*, *elsewhere*) and lock state (*locked* or *unlocked*) was calculated for the average sessions duration per day. There was a highly significant main effect of context, $F(1.46, 1698.55) = 21.67, p < 0.001$. Mauchly's test of sphericity indicated a violation of the assumption of sphericity ($p < 0.05$), so a Greenhouse-Geisser correction was used. Furthermore, there was a highly significant main effect of lock state, $F(1, 1165) = 718.58, p < 0.001$. The interaction between context and lock state was also highly significant, $F(1.67, 1947.04) = 10.82, p < 0.001$. Mauchly's test of sphericity again indicated a violation of the assumption of sphericity ($p < 0.05$), so a Greenhouse-Geisser correction was used. Tuckey post hoc tests revealed significant differences in the average session length at *home* compared to in the *office* ($p < 0.001$) and at *home* compared to at the *other meaningful place* ($p < 0.001$) or *elsewhere* ($p < 0.001$). Furthermore, significant differences in the average session length in *office* compared to *other meaningful place* ($p < 0.019$) were found. There were no significant differences in the average session length between *office* and *elsewhere* ($p = 0.095$) and between *other meaningful place* and *elsewhere* ($p = 0.93$). Figure 10 outlines the corresponding ANOVA descriptives

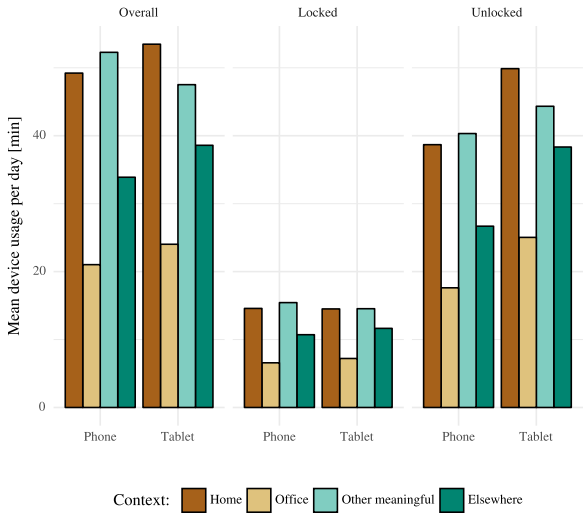

Fig. 14. Grand mean of device usage per day by context

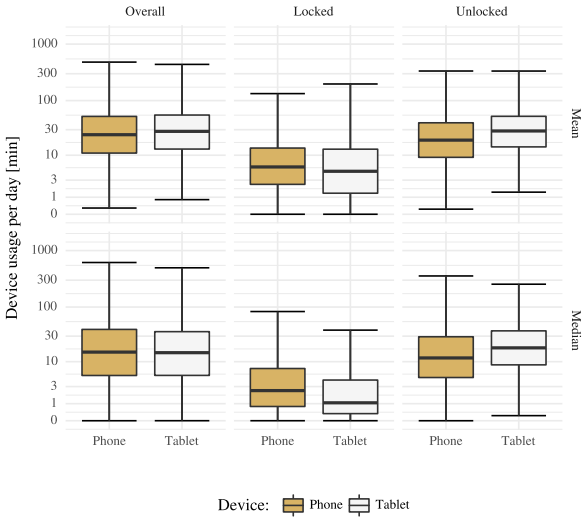

Fig. 15. Distribution of daily usage across devices

4.4 Daily Usage Duration

We found that the average locked device usage per day for phones and tablets is fairly close (36 minutes vs. 25 minutes), as the tablets' longer sessions compensate for the higher number of sessions on phones. Unlocked

usage of tablet devices sums up to 81 minutes per day (median 44 minutes), while phones are used on average 93 minutes per day (median 66 minutes). Overall, phone usage amounts to 126 minutes per day (median 96 minutes) while tablets feature an overall usage of 95 minutes (median 47 minutes). As with individual session length, *home* context accounts for the largest share of usage while *office* has the smallest share per context of daily usage.

The average device usage per day is again dominated by a small amount of devices accumulating an excessive amount of daily usage. Some phones featured an average usage per day of almost 15 hours while the maximum average usage of tablet devices is 7 hours. The median of the overall mean daily usage is, however, 109 minutes for phones and 67 minutes for tablet devices. Figure 15 depicts the distribution of both mean and median of daily usage for locked, unlocked, and overall device usage.

4.5 Diurnal Pattern

The long-term nature of the underlying dataset – some users participate for more than 4.5 years – enables us to analyze diurnal patterns in mobile device usage. For this purpose we measured how much time each user would spend at which days of the week and which hour. Since users participated for quite different periods, each user's usage distribution was scaled to sum up to one. Values across users w ere normalized in the interval [0, 1] with one marking the time frames in which the most mobile device usage occurs while zero implies no device usage at all. Figures 16 and 17 show that diurnal usage patterns appear to be quite different with regards to context, time and day of the week. Mobile devices are most intensely used during weekdays between 09:00 and 17:00 in

Fig. 16. Weekly phone usage by unlocked usage time

Fig. 17. Weekly tablet usage by unlocked usage time.

the office context, e.g., at work or in class. For *home* and *other meaningful*, one can observe different pattern for weekdays and weekends. At *home*, usage during early morning hours seems to be less intense at weekends than during the week while the reverse seems to be true for *other meaningful places*. These patterns are quite obvious for phones while the tablet data seem more noisy. This is at least in part due to the fact that significantly more data was available for phones than for tablets (a total of more than 207 years of active phone usage sessions vs. 12 years of active tablet usage). We note that analyzing diurnal pattern across all users is not ideal, since users with different patterns, e.g., morning person and late riser, blur the overall picture when combined. It is still sufficient to illustrate the point that context, time and day of week seem to have a notable effect on mobile device usage.

When considering day-based statistics, it is also worth pointing out that the extent of device usage tends to vary considerably on the long term. Figure 18 shows the monthly usage of a single phone over the course of 42 months (67,844 usage sessions), with usage varying from 37.5 hours to 9.8 hours per month.

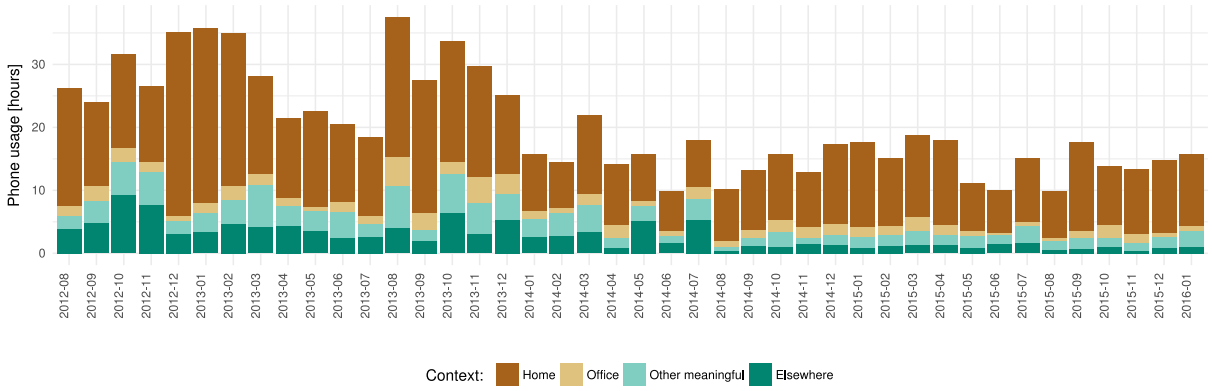

Fig. 18. Usage of a single phone over the course of several months

4.6 Device Unlocking

Apart from analyzing unlocked device usage, we analyzed how users lock their devices based on all devices within the original dataset featuring the required information. Unlocking a device requires either slide-to-unlock or some form of authentication like PIN, password, or graphical pattern. Since the underlying dataset unfortunately only labels graphical pattern-based unlocking explicitly, means of comparing different authentication methods are limited. However, pattern unlock seems to be quite popular, as it is enabled on 35% of the smartphones and on 24% of the tablet devices. Of these devices, 72% are configured to provide visual feedback while entering the pattern, increasing the vulnerability to so-called shoulder surfing attacks, i.e., capturing the secret pattern by looking over the user's shoulder during device unlocking [3, 32]. On 8% of the phones and 15% of the tablets no form of device locking, not even slide-to-unlock, is enabled (see fig. 19).

One aspect of the usability of unlocking mechanisms is the speed at which the device can be unlocked. Using the state machine approach described in fig. 3, we measure the time between turning the device's screen on and unlocking the device, indicated by a *USER_PRESENT* intent broadcasted by the Android system when the device is unlocked. The 20.7 million unlocking sessions we extracted that way, however, also contain sessions in which the user turns the device on but only attempts to unlock it after several minutes (given a long display timeout is configured). We therefore choose an upper limit of 10 seconds and only take shorter unlocking sessions into account, which leaves us with 19.6 million sessions.

A comprehensive real world study conducted by Zezschwitz et al. [34] (n=31) reported that unlocking takes on average 1.5 seconds for PIN-based mechanisms and 3.1 seconds for pattern-based unlocking, concluding "users of the pattern system needed more than twice as much time as PIN users to achieve a successful login". Interestingly, a more recent study by Harbach et al. [7] (n=134) indicates quite the opposite, finding pattern-based unlock to take 0.9 seconds on average while PIN users spend on average 2.0 seconds to unlock their devices.

Our results seem to confirm the observation of Zezschwitz et al. [34] that pattern unlock requires notably more time (especially when looking at the median unlock time) than other unlocking mechanisms like PIN, as the unlocking duration distribution (see fig. 20) illustrates. Looking at sessions shorter than 10 seconds, we find that pattern unlock on smartphones requires on average 2.7 seconds (median 2.3 seconds) while other unlocking methods take 2.5 seconds on average (median 1.8 seconds). On tablet devices, pattern unlock requires on average 3.2 seconds (median 2.6 seconds) while other unlocking methods take 2.9 seconds on average (median 2.2 seconds).

The discrepancy between the observations made by Harbach et al. [7] and Zezschwitz et al. [34] as well as this paper are caused by the way how the unlocking time is measured. While in [34] as well as in our paper, time is counted from the moment the screen is turned on, Harbach et al. [7] does distinguish between *preparation*, which begins after the screen is turned on, and the actual unlock process, which begins by entering the first PIN digit or starting to actually draw the unlock pattern.

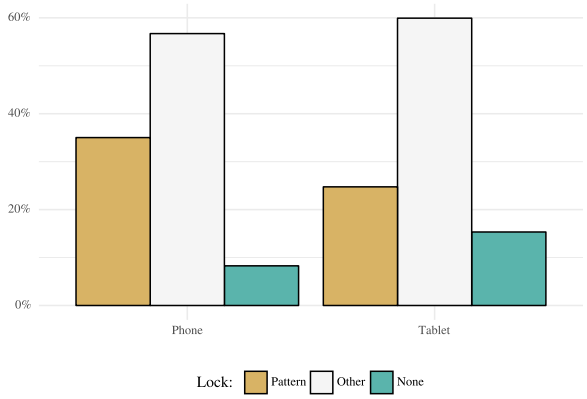

Fig. 19. Usage of different locking mechanisms

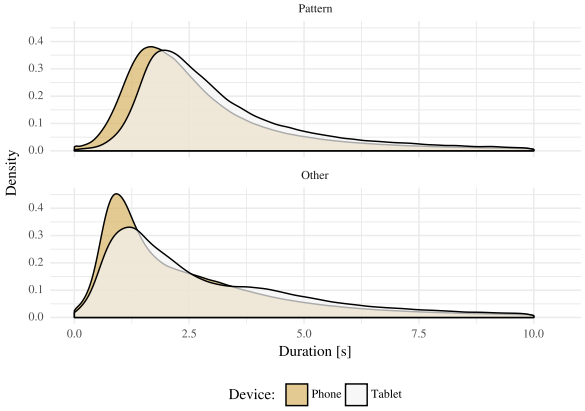

Fig. 20. Density of unlocking session duration

5 LIMITATIONS

The dataset used in this work and consequently the results of our analysis are limited in some ways. First, the dataset only contains logs for Android devices. Results for other mobile platforms might differ. And while the panel of 29,279 participating devices is recruited over the course of several years and a variety of different channels and thus fairly diversified in terms of geographical location, device model, and manufacturer (see figs. 1 and 2), we lack demographic information to make solid statements about how representative the panel actually is.

Another limitation is that usage sessions are not present explicitly in the dataset but have to be derived from secondary features like display power status or phone subsystem events. While we took great care to provide the highest data quality possible, usage session extraction in the end remains an approximation that is to some extent distorted by, e.g., display timeouts. This limitation is universal to all handset-based user studies, since average devices are not (yet) capable of reliably and accurately tracking user attention.

When comparing phone and tablet statistics, it has to be considered that since we do omit days without any user interaction, it would not become apparent if devices are used only infrequently on a large time scale, e.g.,

once per month. While tablets in our analysis are used on average on less days of the observation period compared to phones, we think the discrepancy is not at a scale that prohibits direct comparison. In particular, we used on average 81% of all days recorded on phones and 74% of all days in the observation period recorded on tablets.

The context classification algorithm used in our work is, due to a lack of ground truth, fuzzy, which has to be taken into account when interpreting the results in terms of different contexts.

The dataset does not contain information to infer whether multiple devices, e.g., a phone and a tablet, are owned by the same person. This does limit our results to device-based statistics instead of user-based statistics, which might be more relevant for certain applications (see [5] for such an analysis).

While we argue that the aspects of mobile device usage analyzed in this paper are very useful for a number of fields and applications, there are doubtlessly a number of other device usage features for which a large scale study could provide useful insights, for instance which applications users spend time on or how users switch between different applications. One work that studied those features on a larger scale ($n = 4125$) is Böhmer et al. [1], but given that their data were collected more than 7 years ago and modern smartphones have only been around for about 10 years, a new analysis based on current data would certainly be useful. While the Device Analyzer dataset contains information about installed and used applications, they are obfuscated to protect the privacy of the participants and allow no insight into the purpose of the application or even comparison across users. Therefore, we were unfortunately not able to provide more insight into these aspects within the scope of this study.

6 CONCLUSION

In this work we studied locked and unlocked mobile device usage with respect to device form factor and user context. For our study we extracted a total of 56.3 million usage sessions from 225 billion mobile device usage records using a sophisticated screen power state machine-based approach. By combining anonymized GSM cell IDs, Wi-Fi scan results, and timestamps of records we derived location information for usage sessions. Through making (presumably) reasonable assumptions about standard users' diurnal patterns, we were able to make fair guesses about users' locational context, identifying *home* context for 88 % and *office* context for 80 % of the smartphone devices.

Consistent with previous studies we found high diversity in device usage characteristics, both across sessions and users, varying with more than an order of magnitude. We observed that on average, smartphones are used almost thrice as much per day as tablet devices (60 times vs. 23 times). However, devices are unlocked in only 46 % of the interactions. Given the limited forms of interaction available in locked state, the high share of locked usage indicates that the majority of usage constitutes some form of short information checking. Our results show that 19 % of smartphone usage occurs in *office* context and 27 % in *home* context. Contrary to the number of interactions, we found that the duration of usage sessions is in general more than twice as long for tablets compared to smartphones: on average, unlocked sessions on phones last 307 seconds while tablet usage sessions account for 963 seconds. Thus, the daily usage of both smartphones and tablets are not far off (93 minutes vs. 81 minutes). Again, home context accounts for the largest share of usage while office has the smallest share per context of daily usage.

Our work shows that despite offering similar technical capabilities, smartphones and tablets are used quite differently. While substantial research has been conducted with respect to smartphone usage, little work has been done to analyze tablet usage. With the increasing ubiquity of mobile devices, people tend to simultaneously own and use several devices of different form factors like phones, tablets, and smartwatches. Further research is needed, e.g., on when and why users change between different device types, and to include the newer form factors like smartwatches, -glasses, etc.

The results of our work are applicable to a number of research topics but are also relevant for practitioners. Application developers for instance should make use of the – somewhat surprising – fact that a signification

portion of device interaction already does not involve unlocking and should consider making application parts that are not security sensitive (e.g., switching to another song on a media player) accessible without unlocking. Security researchers could consider applying different security settings based on the location context in order to reduce the perceived burden of user authentication, for instance at home [11]. For behavioral psychologists, our findings provide a comprehensive reference for analyzing compulsive behavior, technostress and smartphone addiction [15, 16, 18], where interaction frequency and duration are key metrics.

Table 3. Comparison of common usage session characteristics for different mobile device usage studies.

	Devices		**Daily interactions**						**Session length [sec]**						**Daily usage [min]**					
	quantity	days	overall		locked		unlocked		overall		locked		unlocked		overall		locked		unlocked	
SMARTPHONES		MEAN	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.
Falaki et al. [4]	255	106	10-250	-	-	-	-	-	10-250	-	-	-	-	-	30-500	-	-	-	-	-
Oliver [21]	17,300	17	87	76	-	-	-	-	68	20	-	-	-	-	101	79	-	-	-	-
Soikkeli [25]	140	21-91	-	-	-	-	20	-	-	-	-	-	207	45	-	-	-	-	73	-
Böhmer et al. [1]	4,125	127	-	-	-	-	-	-	-	-	-	-	-	-	59	-	-	-	-	-
Truong et al. [27]	10	14-20	-	-	-	-	5-105	-	-	-	-	-	-	-	-	-	-	-	-	-
Finley and Soikkeli [5]	561	-	-	-	-	-	-	-	-	-	-	-	245	56	-	-	-	-	-	-
Harbach et al. [7]	134	30	70	57	-	-	40	32	-	-	73	39	355	260	-	-	-	-	-	-
Hintze et al. [10]	1,487	-	58	44	37	24	25	19	165	30	94	11	299	74	117	82	43	18	86	58
Our study	9,861	144	60	48	34	24	27	21	206	147	107	57	307	223	126	96	36	15	93	66
TABLETS		MEAN	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.
Finley and Soikkeli [5]	65	-	-	-	-	-	-	-	-	-	-	-	506	114	-	-	-	-	-	-
Hintze et al. [10]	98	-	27	12	17	6	11	6	414	73	206	15	694	197	112	67	36	7	88	53
Our study	672	230	23	8	15	4	10	4	616	366	271	84	963	656	95	47	25	3	81	44
SMARTPHONES & TABLETS		MEAN	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.	MEAN	MED.
Wagner et al. [31]	16,000	43	57	-	-	-	-	-	116	-	-	-	-	-	123	79	-	-	-	-

A SUPPLEMENTARY MATERIALS

The Device Analyzer dataset consists of logs of key/value pairs of 263 different features – ranging from airplane mode settings to wifi scan events – of which we used only about 10%. The following is a verbatim excerpt from the dataset documentation[6] describing the subset of the available features used in this work. For a comprehensive description of the complete dataset, see [30, 31] or visit https://deviceanalyzer.cl.cam.ac.uk.

- **hf**: Contains information that is frequently collected while the screen is on.
  - **locked**: Whether or not the keyguard is active. When the screen is successfully unlocked this changes to *false*.
- **pause**: Indicates a privacy pause or a manual resume from one. The value *resume* indicates a manual resume. Any number indicates the duration in ms that the pause will be for. Afterwards logging will resume normally.
- **phone**: Contains information about the state of the telephony subsystem.
  - **celllocation**: Contains information about the current network cell.
    - **cid**: The Cell ID for GSM signals.

[6] http://deviceanalyzer.cl.cam.ac.uk/keyValuePairs.htm

        * **lac**: The Location Area Code for GSM signals.
        * **basestationid**: The base station ID for CDMA signals.
        * **networkid**: The network ID for CDMA signals.
        * **systemid**: The system ID for CDMA signals.
    - **ringing**: The phone state changed: There is a new incoming call. (anonymize)
    - **calling**: The phone state changed: There is a new outgoing call. (anonymize)
    - **offhook**: The phone state changed: A connection has been established and there is now an active call. The *value* is always an empty string.
    - **idle**: The phone state changed: The telephony subsystem is idle now.
    - **keyguardremoved**: The key guard was removed (corresponds to `ACTION_USER_PRESENT`). The device can now be operated by a user. Value is always empty. This key is present since version 1.1.5.
- **power**: Information about the battery and whether or not the device is charging.
- **screen**: Contains information about the power state of the display.
    - **power**: This entry is fired whenever the display is turned on or off.
- **shutdown**: Indicates that the device is powering down. Note that this message may or may not be present when the devices is turned off and that you should not rely on it being present, e.g. in case of power failure.
- **startup**: Indicates that all sources are being initialized.
- **system**: Contains information about the state of the system and our software preferences. These keys are stored with every boot.
    - **apiversion**: The highest API version available on the device, e.g. "7" for Android 2.1
    - **device**: A string identifying the hardware of the device, e.g. "hero". Output of `Build.DEVICE`
    - **display**: Information about the device's display. Present since version 1.1.2.
        * **density**: The density that is used to calculate the size of on-screen elements, as a factor of the "default" 160dpi screen, e.g. 1.5 for a 240dpi screen. Note that these numbers only roughly correlate with physical screen density.
        * **dpi**: The display's physical density in x and y dimension, as reported by `DisplayMetrics.xdpi` and `ydpi`. E.g. 254.0x254.0
        * **resolution**: The display's resolution in pixels as reported by the OS, e.g. 480x800
    - **locale**: Represents language code, country code and variant, separated by underscores. Missing values are omitted. Examples are "en", "en_US", "_US", "en__POSIX", "en_US_POSIX". As of version 1.2.0 this key now has two sub-keys. Before, this key contained the user's preferred locale, as returned by `Locale.getDefault()`.
        * **default**: The user's preferred locale, as returned by `Locale.getDefault()`.
        * **current**: The user's currently active locale.
    - **manufacturer**: A string identifying the hardware manufacturer. Output of `Build.MANUFACTURER`.
    - **model**: A string identifying the manufacturer's name for the device, e.g. "Galaxy S2". Output of `Build.MODEL`
    - **settings**: Contains general system settings. Present since version 1.1.2.
        * **nonmarketapps**: Whether non-market apps can be installed. May not be present if not set.
        * **lock**: Whether the lock pattern is enabled. May not be present if not set.
        * **locktactile**: Whether the lock pattern gives tactile feedback. May not be present if not set.
        * **lockvisible**: Whether the Lock pattern is visible. May not be present if not set.
        * **screenoff**: How long the screen stays active without user input before turning off.
    - **swbuild**: The build number of DeviceAnalzyer.

  - **swversion**: The human-readable version string of DeviceAnalyzer.
- **tethering**: Generated when the system's tethering state changes
  - **active**: Lists all *active* tethering devices. If none are present, this entry is omitted.
  - **available**: Lists all *available* tethering devices. If none are present, this entry is omitted.
  - **errored**: Lists all *failed* tethering devices. If none are present, this entry is omitted.
- **wifi**: Contains information about currently visible wifi networks.
  - **scancomplete**: Marker that indicates that a wifi scan finished. The value contains the number of visible APs, which will follow immediately after this marker.
  - **scan**: Results of a scan for wifi networks in range
    * **[BSSID]**: The access point's MAC address. (anonymize)
    * **ssid**: The network name that is displayed to the user. Multiple access points can belong to one network. (anonymize)

ACKNOWLEDGMENTS

We thank the University of Cambridge for providing access to the Device Analyzer dataset and in particular Andrew Rice for instantaneously answering all our questions. Furthermore, we thank the four anonymous reviewers whose constructive criticism and suggestions helped improve and clarify an earlier draft of this paper. Additionally, we thank Klaus-Dieter Labahn for his thorough proofreading. Finally, we gratefully acknowledge funding by the German Federal Ministry of Education and Research as well as *u'smile*, the Josef Ressel Center for User-Friendly Secure Mobile Environments, funded by the Christian Doppler Gesellschaft, A1 Telekom Austria AG, Drei-Banken-EDV GmbH, LG Nexera Business Solutions AG, NXP Semiconductors Austria GmbH, and Österreichische Staatsdruckerei GmbH.

REFERENCES

[1] 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. *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services - MobileHCI '11* January (2011), 47. DOI: https://doi.org/10.1145/2037373.2037383

[2] Guanling Chen and David Kotz. 2000. *A Survey of Context-Aware Mobile Computing Research.* Technical Report.

[3] Alexander De Luca, Marian Harbach, Emanuel von Zezschwitz, Max-Emanuel Maurer, Bernhard Ewald Slawik, Heinrich Hussmann, and Matthew Smith. 2014. Now You See Me , Now You Don't - Protecting Smartphone Authentication from Shoulder Surfers. *Sigchi* (2014), 2937–2946. DOI: https://doi.org/10.1145/2556288.2557097

[4] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. Diversity in Smartphone Usage. *Proceedings of the 8th international conference on Mobile systems, applications, and services - MobiSys '10* (2010), 179. DOI: https://doi.org/10.1145/1814433.1814453

[5] Benjamin Finley and Tapio Soikkeli. 2016. Multidevice mobile sessions: A first look. *Pervasive and Mobile Computing* (in press) (2016). DOI: https://doi.org/10.1016/j.pmcj.2016.11.001

[6] Preetinder S Gill, Ashwini Kamath, and Tejkaran Singh Gill. 2012. Distraction: an assessment of smartphone usage in health care work settings. *Risk Manag Healthc Policy* 5, 1 (2012), 105–14.

[7] Marian Harbach, Alexander De Luca, and Serge Egelman. 2016. The Anatomy of Smartphone Unlocking: A Field Study of Android Lock Screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*.

[8] Daniel Hintze, Rainhard Dieter Findling, Muhammad Muaaz, Eckhard Koch, and René Mayrhofer. 2015. CORMORANT: Towards Continuous Risk-Aware Multi-Modal Cross-Device Authentication. *UbiComp 2015 Adjunct Publication* (2015).

[9] Daniel Hintze, Rainhard D Findling, Muhammad Muaaz, Sebastian Scholz, and René Mayrhofer. 2014. Diversity in Locked and Unlocked Mobile Device Usage. In *UbiComp 2014 Adjunct Publication*. 379–384.

[10] Daniel Hintze, Rainhard Dieter Findling, Sebastian Scholz, and René Mayrhofer. 2014. Mobile Device Usage Characteristics: The Effect of Context and Form Factor on Locked and Unlocked Usage. In *Proceedings of MoMM 2014*.

[11] Daniel Hintze, Muhammad Muaaz, Rainhard Dieter Findling, Sebastian Scholz, Eckhard Koch, and René Mayrhofer. 2015. Confidence and Risk Estimation Plugins for Multi-Modal Authentication on Mobile Devices using CORMORANT. In *Proceedings of MoMM 2015*. 384–388.

[12] Daniel Hintze and Andrew Rice. 2016. Picky: Efficient and Reproducible Sharing of Large Datasets Using Merkle-Trees. *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)* (2016), 30–38. DOI: https://doi.org/10.1109/MASCOTS.2016.25

[13] Borja Jiménez. 2008. *Modeling of Mobile End-User Context.* Ph.D. Dissertation. Helsinki University of Technology.

[14] Jeffrey H Kuznekoff and Scott Titsworth. 2013. The impact of mobile phone usage on student learning. *Communication Education* 62, 3 (2013), 233–252.

[15] Min Kwon, Joon-Yeop Lee, Wang-Youn Won, Jae-Woo Park, Jung-Ah Min, Changtae Hahn, Xinyu Gu, Ji-Hye Choi, and Dai-Jin Kim. 2013. Development and validation of a smartphone addiction scale (SAS). *PloS one* 8, 2 (2013), e56936.

[16] Heyoung Lee, Heejune Ahn, Samwook Choi, and Wanbok Choi. 2014. The SAMS: Smartphone Addiction Management System and Verification. *Journal of Medical Systems* 38, 1 (2014), 1. DOI: https://doi.org/10.1007/s10916-013-0001-1

[17] Yu-Kang Lee, Chun-Tuan Chang, You Lin, and Zhao-Hong Cheng. 2014. The dark side of smartphone usage: Psychological traits, compulsive behavior and technostress. *Computers in Human Behavior* 31 (2014), 373–383. DOI: https://doi.org/10.1016/j.chb.2013.10.047

[18] Yu-Hsuan Lin, Yu-Cheng Lin, Yang-Han Lee, Po-Hsien Lin, Sheng-Hsuan Lin, Li-Ren Chang, Hsien-Wei Tseng, Liang-Yu Yen, Cheryl C H Yang, and Terry B J Kuo. 2015. Time distortion associated with smartphone addiction: Identifying smartphone addiction via a mobile application (App). *Journal of Psychiatric Research* 65 (2015), 139–145. DOI: https://doi.org/10.1016/j.jpsychires.2015.04.003

[19] Gonçalo M. S. Marques and Rui Pitarma. 2016. Smartphone Application for Enhanced Indoor Health Environments. *Journal of Information Systems Engineering & Management* 1, 4 (2016), 1–9. DOI: https://doi.org/10.20897/lectito.201649

[20] Hendrik Müller, Jennifer Gove, and John Webb. 2012. Understanding Tablet Use - A Multi-Method Exploration. *Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI'12)* (2012), 1–10. DOI: https://doi.org/10.1145/2371574.2371576

[21] Earl Oliver. 2010. The Challenges in Large-Scale Smartphone User Studies. *Proceedings of the 2nd ACM International Workshop on Hot Topics in Planet-scale Measurement - HotPlanet '10* (2010), 1. DOI: https://doi.org/10.1145/1834616.1834623

[22] Antti Oulasvirta, Tye Rattenbury, Lingyi Ma, and Eeva Raita. 2011. Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing* 16, 1 (jun 2011), 105–114. DOI: https://doi.org/10.1007/s00779-011-0412-2

[23] Charlie Pinder, Russell Beale, and Robert J Hendley. 2016. Accept the Banana : Exploring Incidental Cognitive Bias Modification Techniques on Smartphones. *CHI Extended Abstracts on Human Factors in Computing Systems* (2016), 2923–2931. DOI: https://doi.org/10.1145/2851581.2892453

[24] Husnjak Siniša, Peraković Dragan, and Cvitić Ivan. 2016. Relevant Affect Factors of Smartphone Mobile Data Traffic. *Promet - Traffic&Transportation* 28, 4 (2016), 435–444. DOI: https://doi.org/10.7307/ptt.v28i4.2091

[25] Tapio Soikkeli. 2011. *The effect of context on smartphone usage sessions.* Master's Thesis. Aalto University School of Science.

[26] T. Soikkeli, J. Karikoski, and H. Hammainen. 2011. Diversity and End User Context in Smartphone Usage Sessions. *2011 Fifth International Conference on Next Generation Mobile Applications, Services and Technologies* (Sept 2011), 7–12. DOI: https://doi.org/10.1109/NGMAST.2011.12

[27] Khai N. Truong, Thariq Shihipar, and Daniel J. Wigdor. 2014. Slide to X: Unlocking the Potential of Smartphone Unlocking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14).* ACM, New York, NY, USA, 3635–3644. DOI: https://doi.org/10.1145/2556288.2557044

[28] Niels van Berkel, Chu Luo, Theodoros Anagnostopoulos, Denzil Ferreira, Jorge Goncalves, Simo Hosio, and Vassilis Kostakos. 2016. A Systematic Assessment of Smartphone Usage Gaps. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 4711–4721. DOI: https://doi.org/10.1145/2858036.2858348

[29] Hannu Verkasalo. 2008. Contextual patterns in mobile service usage. *Personal and Ubiquitous Computing* 13, 5 (mar 2008), 331–342. DOI: https://doi.org/10.1007/s00779-008-0197-0

[30] Daniel T. Wagner, Andrew Rice, and Alastair R. Beresford. 2013. Device Analyzer: Large-scale mobile data collection. In *Big Data Analytics workshop, ACM Sigmetrics 2013.*

[31] Daniel T. Wagner, Andrew Rice, and Alastair R. Beresford. 2013. Device Analyzer: Understanding smartphone usage. In *10th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services.*

[32] Susan Wiedenbeck, Jim Waters, Leonardo Sobrado, and Jean-Camille Birget. 2006. Design and evaluation of a shoulder-surfing resistant graphical password scheme. *Proceedings of the working conference on Advanced visual interfaces - AVI '06* (2006), 177. DOI: https://doi.org/10.1145/1133265.1133303

[33] Yafei Yang, Lu Xiao, Yongjin Kim, and David Julian. 2009. Case Study: Trust Establishment in Personal Area Networks. *Proceedings of ISWPC 2009* (2009), 1–5.

[34] Emanuel Von Zezschwitz, Paul Dunphy, and Alexander De Luca. 2013. Patterns in the Wild: A field study of the usability of pattern and pin-based authentication on Mobile Devices. *Proceedings of the 15th International Conference on Human-Computer Interaction with Mobile Devices and Services* (2013), 261–270.

Received February 2017; revised April 2017