Advancing the Understanding of Android Unlocking and Usage

by

Lina Qiu

BEng in Computer Science and Engineering, University of Electronic Science and Technology of China, 2015

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Applied Science

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Electrical and Computer Engineering)

The University of British Columbia
(Vancouver)

May 2018

© Lina Qiu, 2018

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, a thesis/dissertation entitled:

ADVANCING THE UNDERSTANDING OF ANDROID UNLOCKING AND USAGE

submitted by LINA QIU in partial fulfillment of the requirements for

the degree of MASTER OF APPLIED SCIENCE

in ELECTRICAL AND COMPUTER ENGINEERING

Examining Committee:

KONSTANTIN BEZNOSOV, ELECTRICAL AND COMPUTER ENGINEERING
Co-supervisor

ALEXANDER DE LUCA, RESEARCH SCIENTIST OF GOOGLE, ZURICH
Co-supervisor

Supervisory Committee Member

JULIA RUBIN, ELECTRICAL AND COMPUTER ENGINEERING
Additional Examiner

Additional Supervisory Committee Members:

Supervisory Committee Member

Supervisory Committee Member

Abstract

Research efforts have been made towards creating mobile authentication systems to better serve users' concerns regarding usability and security. While previous works have revealed real world smartphone authentication usage patterns, several aspects still need to be explored. In this research, we fill some of these knowledge gaps, including how age influences smartphone use. To this end, we performed a two-month long field study on a diverse North American study pool (N = 137). We examined how smartphone usage correlates with users' ages, their choice of unlocking mechanisms (e.g., PIN vs. Pattern) and the types of activities they undertook while unlocking their phones. Study results reveal that there are indeed significant differences across age and unlocking mechanisms. For instance, older participants interacted significantly less-frequently with their devices, and for a significantly shorter amount of time each day. Fingerprint users had significantly more device sessions than other mechanism groups. In addition, we also observed that most participants regularly shared their devices with others, while they also likely underestimated the sensitivity of the data stored on them. Overall, these observations provide important messages for designers and developers of smartphone authentication systems.

Lay Summary

This thesis presents the results of a research project to create a better understanding of real world smartphone authentication usage patterns, including how age correlates with smartphone use. To conduct this research, we performed a two-month field study on a diverse North American study pool (N = 137). We examined how smartphone usage correlates with users' age, their choices of unlocking mechanisms (e.g., PIN vs. Fingerprint) and the types of activities they undertook during the unlocking process. Study results show that there are indeed significant differences across age and unlocking mechanisms. For instance, older age groups interacted with their devices significantly less frequently, and for a significantly shorter amount of time each day. Fingerprint users had significantly more device sessions than PIN and Swipe/None users. In addition, we also observed that most participants regularly shared their devices with others, while they also likely underestimated the sensitivity of the data stored on their devices.

Preface

This research was the product of a fruitful collaboration between the author of the thesis and the following people: Ildar Muslukhov and Konstantin Beznosov (supervisor) from the University of British Columbia (UBC), and Alexander De Luca (co-supervisor) from Google, Zurich. The idea of the project was built upon the following previous work:

A. Mahfouz, I. Muslukhov, and K. Beznosov. Android users in the wild: Their authentication and usage behavior. *Special Issue on Mobile Security, Privacy and Forensics*, Volume 32, Pages 50-61, Elsevier, October 2016.

The work presented herein is prepared to submit to CHI Conference on Human Factors in Computing Systems, 2019 and is in preparation:

L. Qiu, A. De Luca, I. Muslukhov, and K. Beznosov. Lock of Ages: Towards Understanding Age Influence on Smartphone Authentication. *In preparation*.

I was responsible for designing and implementing the study application (app), recruiting participants, collecting and analyzing the data, and writing the manuscript, while Ildar Muslukhov helped set up the backend data collection server. All co-authors actively participated in discussions on the app design, participant recruitment, data collection and analysis, and paper writing process. For this study, I obtained ethics approval from the Behavioural Research Ethics Board (BREB) at UBC. Approval H16-00343, titled “Android Usage”.

Table of Contents

List of Tables

List of Figures

Acknowledgments

First and foremost, I would like to thank my supervisor, Konstantin Beznosov, and co-supervisor, Alexander De Luca, for giving me this opportunity to work on this very interesting project; I am grateful for their patient guidance and support throughout this journey.

Next, I would like to thank my collaborator, colleague, and supportive friend, Ildar Muslukhov, for providing me with a great deal of help and many suggestions, both on the technical and intellectual level, during the various stages of this research.

Further, I would like to thank an additional colleague and supportive friend, Primal Wijesekera, for his feedback and constructive discussions on how to write my thesis, and his detailed review of it. I would like to thank Ross Sheppard for his invaluable help with proofreading the conference paper upon which this thesis is based.

I would also like to thank Ahmed Mahfouz for performing his initial project on understanding Android users in the wild, which inspired the idea for the current project. Last but not least, I would like to thank my beloved family, especially my parents, sisters and brother, and all those who patiently supported me during this journey.

Dedication

To my beloved parents, siblings,
and all my family members who
never ceased from supporting me.

Chapter 1

Introduction

Smartphones have become one of today's most commonly used computing platforms, with the Android being one of the world's most popular mobile operating systems [16]. Advances in these devices' capabilities enable the storage of (and access to) large amounts of data, some of which can be retained as sensitive or private in nature [15]. Also, due to the small size and high mobility of such devices, unauthorized access to sensitive data has become a significant threat. For instance, it has been shown that one out of every five users in the United States has accessed another user's smartphone without permission [14].

To protect smartphones from unauthorized access, all mobile operating systems provide authentication-based device locking. However, more than 40% of all smartphone users do not use a secure locking mechanism [5]. Furthermore, most smartphone users tend to choose easy-to-guess unlocking secrets [3]. The inconvenience of currently deployed unlocking methods, lack of motivation, and lack of awareness about the sensitivity of the data stored on their devices are often used to justify why a secure lock is not being used [5, 9, 15].

Recent studies [7, 8, 12, 13], some of which were conducted in the wild, have shed light on how and under which circumstances smartphone users employ different unlocking mechanisms. For instance, Mahfouz et al. [13] used a student sample to investigate authentication process parameters, such as the time required

to unlock devices, authentication error rates, and the types of apps used within each session. Harbach et al. [7] focused on authentication speed, error counts, and the types of errors observed in different unlocking mechanisms, based on a study sample of 134 participants. While these efforts have provided the first valuable insight on the performance of smartphone unlocking mechanisms in the wild, there are still plenty of unknowns.

Such examples include the unknown correlations between users' (un)locking behaviours and factors like age and the undertaken activity. Specifically, a recent online survey on secure smartphone locking across eight countries suggested that age might be linked to the locking behaviour of users, e.g., older users were significantly less likely to use a secure lock screen [8]. It is, however, still unclear how age affects smartphone unlocking due to the lack of real-world studies focusing on age. In addition, it is also unknown if unlocking statistics correlate with the type of activity a user undertakes while unlocking their device.

In order to fill this knowledge gap and provide initial insights into how age correlates with smartphone authentication, we conducted a longitudinal field study with 137 participants. The main goal was to show what behavioural patterns in authentication were influenced by age and which were similar across age groups. In contrast to previously published research, our participants had diverse backgrounds and covered many age groups. Each participant was required to install our smartphone usage monitoring application and run it for at least two months. The monitoring app collected detailed data on (un)locking, session usages, activities, etc. We focused our analysis on studying how user age, the used unlocking mechanism (and both together), and current activity correlated with (un)locking behaviour.

The results of the analysis revealed that age has a statistically significant impact on the average number of sessions and the average daily usage length. Participants of older ages tended to interact with their phone less frequently, and for a significantly shorter period of time per day. These variables are important since this means that, for instance, authentication takes up a greater portion of

users' overall device usage. For other patterns such as error rates, the results were rather consistent across all age groups. While one might hypothesize that younger users perform better at (un)locking, our results do not support this. This means that, for existing authentication systems, there is still room for improvement. We also found that the unlocking mechanism chosen significantly correlates with the average number of sessions, which corroborates with findings reported by Harbach et al. [9]. We extended the previous research to incorporate Fingerprint users by showing that they tended to have significantly more sessions than PIN and Swipe/None users. Furthermore, we found that about 50% of the unlocks happened while people were moving. In addition, while over 60% of the participants shared their device with others, most of them shared with their family members and friends, which corroborates with the iPhone sharing patterns reported by Cherapau et al. [3].

Chapter 2

Related Work

Despite the availability of a variety of smartphone unlocking methods (e.g., PIN, alphanumeric passwords, unlock patterns, and biometrics), the adoption of smartphone locking still falls below the expected rate [9, 18]. A recent study [5] showed that most smartphone users are likely to underestimate the sensitivity of their data and how they access it with their devices. Thus, many users do not protect their devices or data properly [8]. One potential reason for this is the inconvenience caused by unlocking and the time this takes [9]. As a consequence, many alternative systems were proposed to make the authentication process easier [1, 10, 11, 17]. SnapApp [1], which provides a time-constrained quick-access option, is an example that reduces the authentication workload by keeping users logged-in in a more secure way than having their device unlocked all the time. However, this approach allows for only limited improvements on usability and the security of the authentication systems, as such short-term access is limited to ten subsequent uses within ten minutes of the last secure unlock (e.g., PIN, Android Pattern, etc.).

To design new mechanisms that are in line with users' real needs, researchers and designers need data on how users utilize their devices in the real world. While there is little information on whether current authentication systems match users' expectations with respect to usability and security, a few studies have provided insights [7, 8, 12, 13].

Mahfouz et al. [13] studied how different smartphone use patterns correlate with the time it takes users to unlock their device, how often users make a mistake during authentication, and which authentication methods users choose for device locking. The findings suggest that users who lock their devices interact with them more frequently and for longer sessions than those who do not. In addition, the cost of unlocking is low when compared with overall smartphone usage and users do not mind adopting unlocking methods with a higher error rate (e.g., Android pattern), as long as they allow faster input of the unlocking secret.

Hintze et al. [12] investigated the number of interactions performed on smartphones per day, the average interaction duration, and the total daily device usage time through use of a state machine based on screen on/off events. Here the authors analyzed mobile device data logs from 1,960 Android smartphones. These logs were collected by the authors of the Device Analyzer project. The authors report that, on average, participants interacted with their devices 57 times per day, among which 43% of time the device was actually unlocked, and the total daily device usage was 117 minutes.

Harbach et al. [8] conducted a global-scale survey on Google Consumer Surveys (GCS) with 8,286 participants from eight countries to investigate whether users' attitudes towards smartphone unlocking differed between various nationalities. The findings indicate that demographic differences, including both nationality and age, should be considered when designing new authentication systems for smartphones. The authors also conclude that, despite the apparent differences between nationalities, the inconvenience of unlocking is still one of the major reasons for the low adoption rate of current authentication systems, especially for older users. The researchers used an online survey to investigate how age affects users' adoption rates of and attitudes towards secure lock screens. Alternatively, we employed real world data to look at how age correlates with usage, e.g., the number of sessions per day and other parameters.

In another month long field study, Harbach et al. [7] collected data from a subset of PhoneLab users, all of which were affiliated with a university. The

authors instrumented LG Nexus 5 smartphones (with Android 4.4) to study the performance of Android unlocking mechanisms in situ. They found that PIN users take longer to unlock while committing fewer errors than pattern users, who tend to unlock their phones more frequently and are more prone to making errors. However, on average, PIN and pattern users spend a similar amount of total unlocking time. In addition, the authors offer a benchmark against which any newly designed unlocking mechanisms can be evaluated.

While the previous work provides many insights on smartphone authentication in situ, all the aforementioned studies suffer from their samples being skewed towards predominantly tech-savvy young participants. In addition, we are unaware of any prior real world studies that have investigated how various age groups differ in their smartphone unlocking practices. To fill this gap, we have conducted a field study with a large and diverse participant pool, focusing our analysis on how smartphone authentication behaviour correlates with participants' ages.

Chapter 3

Research Questions

The main research questions we wanted to address and the rationale behind each such question are as follows:

- **RQ1** - *What influence does age have on smartphone (un)locking?*
  Answering **RQ1** provides a better understanding of whether smartphone users' (un)locking behaviours, e.g., how frequent they unlock their devices, how long they interact with their devices after each unlock, etc., differs with age. This aids the research community in assessing the necessity of designing efficient authentication systems for different age groups.

- **RQ2** - *How does the choice of unlocking mechanisms correlate with the (un)locking behaviour of smartphone users?*
  A previous study [9] revealed that smartphone users' (un)locking behaviour differs across authentication systems. Answering **RQ2** would help improve our understanding of whether these differences in usage hold true for a more diverse study sample, including participants of older age groups, and for new authentication systems, e.g., Fingerprint. We will discuss the diversity of our study sample in detail in Section 5

- **RQ3** - *What is the correlation between the type of user activity and user unlocking behaviour?*

Answering **RQ3** provides a better understanding of how smartphone users react to usage scenarios during different activities, in terms of their unlocking behaviour, including unlocking frequency and error rates, etc. Understanding this question can provide insight into how platforms can support incremental authentication, or different approaches matching the nature of activities users are involved in.

- **RQ4** - *How frequent is device sharing?*
- **RQ5** - *Whom do users share their devices with?*
- **RQ6** - *How does users' preferences for device sharing correlate with age?*
  Answering **RQ4-6** provides insight into smartphone users' device sharing habits, including sharing frequency, targets, and age correlations. This can further assist us in evaluating the to what extent should authentication systems support sharing use cases.
- **RQ7** - *What are the apps that users access the most, and how sensitive are these apps from the point of view of the users?*
- **RQ8** - *How do users' general perceptions of app sensitivity correlate with age?*
  Finally, answering **RQ7** and **RQ8** provides a better understanding of smartphone users' most used applications and their perceptions of the sensitivity of specific apps or all apps in general. In particular, **RQ7** provides a deeper insight into how users perceive the sensitivity of the apps they accessed most, while **RQ8** focuses more on whether users' general perceptions of app sensitivity differ between age groups. Having a better understanding of these two important research questions can help us and the wider research community to further estimate if authentication can be omitted for certain use cases, and whether this can hold true for all age groups.

Chapter 4

Methodology

4.1 Data Collection

Our participants installed a custom-built study application (study app) on their smartphones and ran it for at least 60 days (the first 60 days were used for the analysis). The study app ran in the background and collected relevant usage data. To reach a broad audience, we made the app available through the Google Play Store.

At the first launch, the application presented a consent form to the participant, where they were required to provide consent in order to participate in the study. Once consent was obtained, the application offered participants the opportunity to opt out of providing certain data records. In particular, we made the collection of activity data optional. Afterwards, the app directed the participants to an entry survey that focused on collecting demographic data. In addition, participants were asked to report which unlocking mechanisms they were currently using, if any.

The app required certain types of access permission in order that it be able to monitor important events in Android, such as device unlocking. For instance, users were required to enable the device administrator and usage statistics privileges for the app. Each participant received instructions on how to activate such permission for the study application.

After obtaining consent and receiving the required permission, the app began to collect the data. In particular, the app recorded all lock and unlock events and logged the start and end time stamps of each user session. In addition, the app recorded the names of all apps that the participant opened during each session. The app also collected user activity data, but only if the participant consented to its collection. Finally, the app collected user responses to two surveys. We describe all collected data types in details below.

4.1.1 Locks/Unlocks/Failed Unlocks

To detect smartphone locking/unlocking events, the app logged both when the screen turned off and on and the *keyguard lock/remove* event. To account for failed unlocking attempts, it also monitored *password failure* events. In addition, the app read smartphone settings to differentiate between two different modes of device locking: that is, if the device was locked due to an autolock, i.e., a certain timeout after inactivity, or if the device was locked manually by a user who clicked the power button. The API to read this configuration, however, was only available in Android 6 and older models. Thus, we could not programmatically differentiate as to how the device was locked when participants used Android 7. Thus, participants using Android 7 or newer models were asked to report their configurations in the entry survey. To help them do so, we provided instructions on how to read this value in the settings of the device. Due to similar limitations, our app was only programmatically capable of detecting the used unlocking mechanism for devices running Android 5 or older versions. For participants using Android 6 or later versions, our study app asked them (during the entry and exit surveys) the types of unlocking mechanisms they used. Finally, to detect when a user changed their password for device unlocking, the app monitored the *password change* events. Every time such events were detected, the app asked the participants the changes they had made, i.e., whether they had changed the authentication method, unlocking secret code, or both.

4.1.2 User Sessions

For the purpose of this study, we define a *user session* as the period of time between the user unlocking the screen and then locking again by an explicit locking action or through a timeout. Within each user session the study app recorded all events beginning when the user first interacted with a new application, i.e., the application that the user saw on their screen. When such a change was detected, the app logged the current date and time together with the name of the new foreground application. Note that each application was identified by both (1) the application package name, which is unique throughout Google Play Store, and (2) the application name, which was readable, but not necessarily unique. For example, the application identifier for WhatsApp is *{com.whatsapp;WhatsApp}*.

4.1.3 Activities

The app detected and collected the participant's activities (e.g., running, tilting,[1] walking [6]) as soon as any change in activity type was reported by Android. This activity data was collected only if the participant provided consent. The aim was to understand their session usage patterns and locking behaviour to figure out how to optimize the locking approaches available in the platform.

4.1.4 Device Sharing Survey

The Device Sharing Survey sampled the participants' experience with sharing their smartphone. Presented once a week, the survey asked each participant to recall all instances of them sharing their device during the past week. By choosing a one-week recall, we aimed to reduce the memory burden on the participants and, thus, to make it easier for them to provide the correct information. The survey also asked the participants to categorize the people they shared their devices with; i.e., friends, family, roommates, and others. In addition, participants were also asked whether they had had any concerns that someone from their social circle would

[1] Tilting the device around the horizontal or vertical axes

access their smartphone without their permission. Participants were allowed to postpone answering this survey for either an hour or a day.

We argue that device sharing habits are relevant for authentication. The most important reason for this is that sharing with close people potentially requires authentication systems to be compatible with this behaviour.

4.1.5 Contextual Survey

We incorporated into the study app a contextual survey to measure the participants' perceptions of the sensitivity of their apps in specific usage contexts. Experience sampling has been used in previous research (e.g., see [4, 9]) and has allowed researchers to better understand how a participant's context impacts on their security and privacy decisions.

In particular, participants were asked to quickly assess their surroundings and report if someone were able to view the contents of their smartphone's screen. Afterwards, the participants were asked to rate in that context the sensitivity of the apps they had spent the most time using. Throughout the day, the study app randomly selected unlocking events and presented the sensitivity survey to the participants. In order to keep this task unobtrusive, while allowing coverage of a wide range of contexts, the study app dynamically adjusted the likelihood of presenting the survey so as to present it six times a day at most. Participants were allowed to skip the survey by clicking the "Not now" button.

We argue that participants' perceptions of app sensitivity are also relevant to authentication, as this can help evaluate whether authentication systems are necessary under all use cases.

Both surveys were implemented as mini-questionnaires. The questions can be found in Appendix A.

4.2 Data Transmission

To protect the confidentiality of the collected data, the study app used an app-specific location on the internal storage of the Android OS. Such storage is protected by the Android's access control subsystem, and is only readable for the corresponding app. In addition, we encrypted all data logged throughout the day with a symmetric encryption key, generated at run time. We encrypted this key with a hard-coded public key, and then appended it to the encrypted data logs before submitting them to the back-end server.

Encrypted logs were uploaded to our back-end server once a day, near midnight. Throughout the study, we downloaded the new data on a daily basis. After decrypting the data, we checked for data corruption; if this had not taken place, we added the data to our dataset.

Chapter 5

Participants

We recruited our participants from North America (US and Canada) through Amazon Mechanical Turk, Twitter, Facebook, our university mailing lists, and The Sample Network.[1] Android smartphone users who were 19 years of age or older were eligible to participate in our study. Everyone was allowed to participate using only one device.

Each participant who ran the study app for 60 days received a compensation of $40 USD and was entered into a raffle for one iPad Pro 2. The chances for each participant to win the raffle were additionally increased each time they answered device sharing and contextual surveys. As part of the compensation, we provided a report to each participant with a statistical description of how they used their smartphone during the study.

In order to make sure that the data collection process was robust, we conducted a pilot study for 15 days with six participants. We obtained approval from the research ethics board of our university before conducting the pilot study.

Overall, we recruited 276 participants. Considering that all participants began the study at different times, all data[2] used in the analysis were collected between December 8, 2016 and August 10, 2017. Out of the 276 participants, 185 completed

[1] The Sample Network (http://thesamplenetwork.com/) was used for recruitment in the US only.
[2] A figure that visualizes the data collection process is presented in Appendix B

the study by providing us with their data for 60 days or more. We note that due to a technical issue, our back-end server had been failing silently for 10 days. While the study app did resubmit data for that period for most participants, we were unable to recover the data from that 10-day period for 43 of our participants, resulting in usable data from a total of only 142 participants.

For each participant, we selected only the first 60 days of their usage. To make sure that we did not mix data from different authentication methods, we also analyzed how often our participants changed their unlocking methods. If a participant changed their unlocking method during our 60-day study period, we verified whether one of these methods accounted for 95% or more of the log records. If so, we retained the participant's data. Otherwise, we removed the data from the analysis; we did this for three such participants.

Finally, we excluded two participants who specified that their unlocking mechanism was "knock (pattern)" and "hold for a set amount of time". This reduced our participant pool down to 137 total participants.

Of the 137 participants whose data we included in the analysis, 123 (89.8%) were from the US, 81 (59.1%) were female and 56 (40.9%) were male. Ages ranged from 19 to 63 years, with a mean age of 40 and median age of 38 (*SD* = 12.5). Participants had diverse education levels, with 56 (41%) having a high school diploma and 35 (26%) having earned a Bachelor's degree. The occupations and salaries of our participants also varied, as shown in Table 5.1.

To evaluate our study sample, we compared it against the smartphone ownership population in the US reported by the Pew Research Center [2]. Statistical results did not reveal any significant differences between our participants' demographics and the one presented in the Pew Research Center report in terms of age ($\chi^2 = 20$, $p = .22$), gender ($\chi^2 = 2$, $p = .16$), education level ($\chi^2 = 15$, $p = .24$), or salary ($\chi^2 = 20$, $p = .22$). We divided all participants into five age groups, based on age group definitions from the Pew Research Center report. The distribution of the participants across all five age groups is shown in Table 5.1.

Parameter	Property	# of participants
Residence	US	123
	Canada	14
Gender	Female	81
	Male	56
Age	19-24	17
	25-34	42
	35-44	26
	45-54	31
	55-63	21
Education	Less than High School	1
	High School	56
	Professional School	23
	University (Bachelor's)	35
	Master or PhD	17
	Other	5
Occupation	Management	9
	Professional	27
	Clerical Support Worker	16
	Service and Sales Worker	18
	Craft and Trades Worker	6
	Machine Operator	1
	Elementary Occupation	3
	Student	14
	Self-employed	3
	Unemployed/Retired/Disabled	28
Salary	Less than $30,000	37
(US, N=123)	$30,000-$49,999	17
	$50,000-$74,999	34

	$75,000-$99,999	17
	$100,000+	17
	Prefer not to specify	1
Salary	Less than $30,000	5
(Canada, N=14)	$30,000-$49,999	4
	$50,000-$74,999	1
	$75,000-$99,999	0
	$100,000+	0
	Prefer not to specify	4

Table 5.1: Participant demographics, $N = 137$

Chapter 6

Results

6.1 (Un)locking Behaviour and Age

Figure 6.1 outlines the distribution of unlocking mechanisms for participants among the different age groups. We conducted statistical tests on usage data, including session lengths, the number of sessions, daily usage lengths, error rates, and the number of auto/manual locks among the predefined age groups and in terms of different unlocking mechanisms. We consider all of these factors relevant for authentication behaviour as they are potentially influenced by or influence the (choice of) authentication method. For example, if a group has a great number of sessions per day, they are exposed to authentication more often.

As most significance tests require individual data items to be independent, we therefore averaged the usage data per participant before conducting the tests. As shown in Table 6.1, we found that usage data including number of sessions, daily usage lengths, auto/manual locks were significantly affected by age, whereas others were not. This emphasizes that designers and developers should take age into consideration while designing new authentication systems in smartphones. To investigate unlocking differences, we conducted the significance tests among all mechanism groups, including “Password”, and we did not find any significant differences in all analyzed usages between “Password” and other mechanism

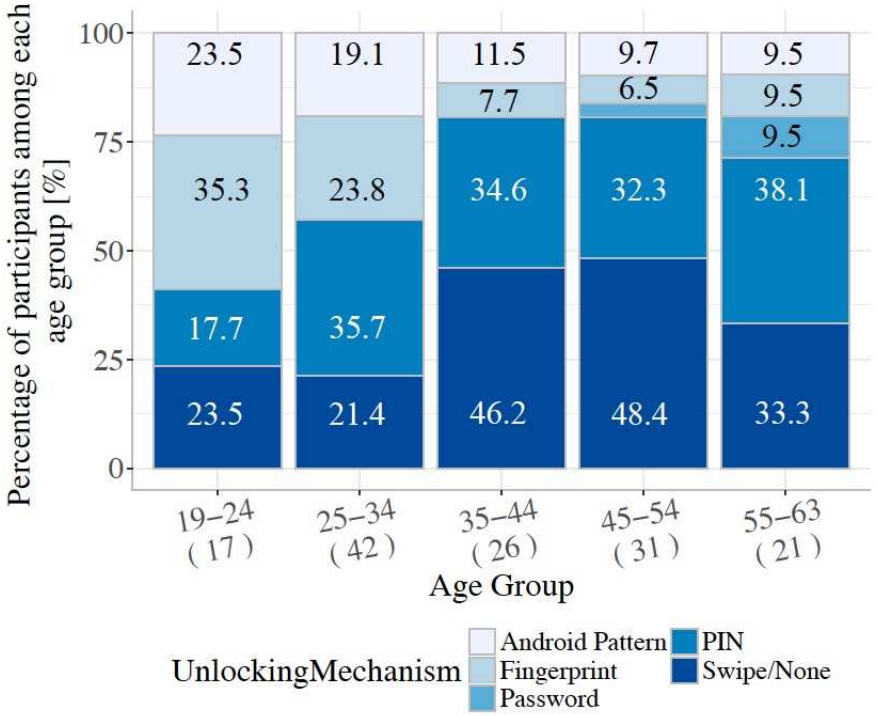

Figure 6.1: Unlocking mechanism distribution among age groups. Numbers below each age group represent the total number of participants from that group.

Usage statistic	*Age Matters?*
Session Lengths	No
Number of Sessions Per Day	Yes
Daily Usage Lengths	Yes
Error Rates	No
Auto/Manual Locks	Yes
Whether authentication happened at still / on the move	Yes
Unlocking Mechanism (at still / on the move)	Yes

Table 6.1: Overview of the usage patterns that were and were not significantly influenced by age.

groups. Thus, we dropped the three password participants while presenting the results for unlocking differences, considering that the size of "Password" group (3) was very small and it lacks sufficient measurement power to perform the tests. In the following subsections, we explain the age and unlocking differences for each dimension of usage separately in detail.

6.1.1 Session Lengths

In total, our data set contained 260,735 user sessions. On average, each session lasted 11.51 minutes (SD = 15.46 minutes). Figures 6.2 and 6.3 show the distribution of average session length across various age groups and unlocking mechanisms, respectively. There was no statistically significant difference in the length of the sessions, either across all age groups (Kruskal-Wallis $\chi^2 = 9.41$, $p = .052$) or across all unlocking mechanism groups (Kruskal-Wallis $\chi^2 = 3.84$, $p = .43$).

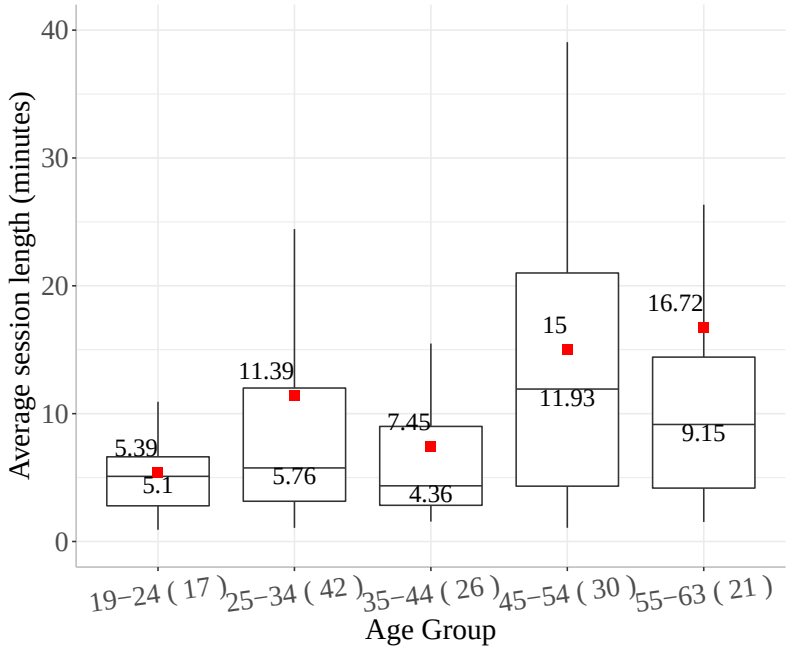

Figure 6.2: Average session length per age group (minutes), $N = 136$. The red squares represent the means.

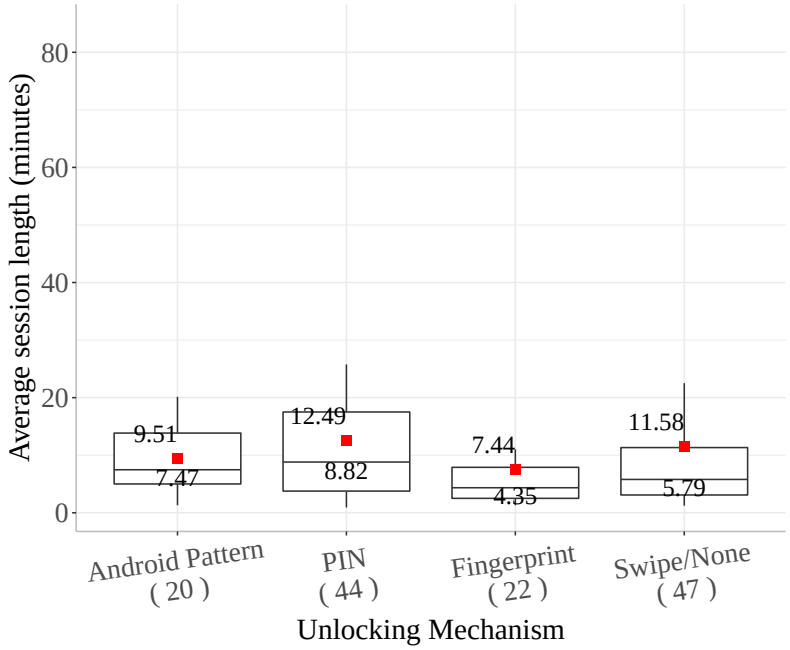

Figure 6.3: Average session length per unlocking mechanism (minutes), $N = 133$ *(the three password participants were removed from the figure, as the sample size was too small)*. The red squares represent the means.

6.1.2 Number of Sessions Per Day

On average, participants had 32 sessions daily with their smartphones ($SD = 26$). Figure 6.4 shows that younger participants interacted more frequently with their devices than older participants (Kruskal-Wallis $\chi^2 = 27.98$, $p < .001$). Further analysis (Bonferroni-corrected Conover-Iman post-hoc) revealed that participants in the "55-63" group had significantly fewer sessions than participants in the younger groups (except for the "45-54" age group). We also found that participants from the age group "45-54" interacted with their devices significantly less than those from the "19-24" and "25-34" groups. The means and standard deviations for each group are presented in Table 6.2.

In addition to the age group correlation, we also tested the impact of the type of the unlocking mechanism on the number of sessions, which turned out to be

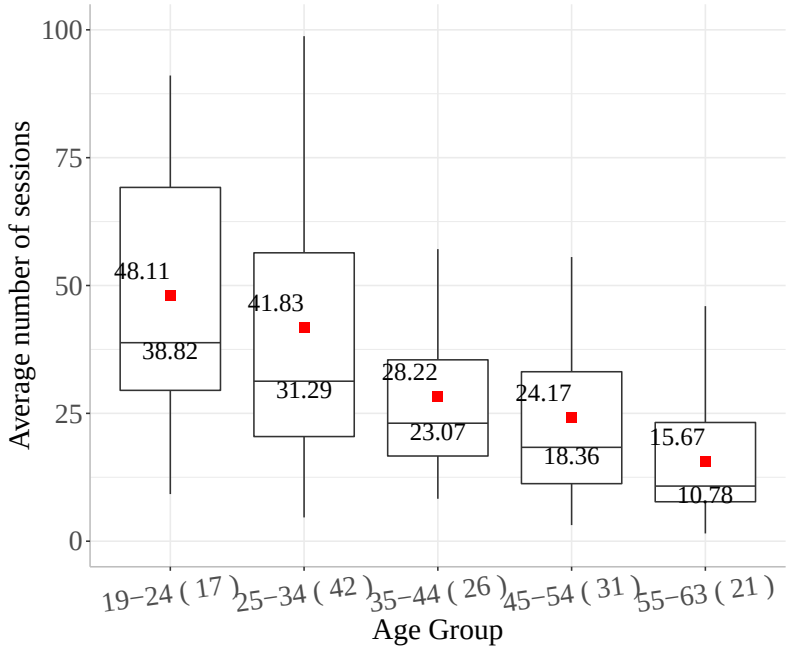

Figure 6.4: Average number of sessions per age group, $N = 137$. The red squares represent the means.

Age group	*Mean*	*Std*
19-24	48.11	23.78
25-34	41.83	33.03
35-44	28.22	17.22
45-54	24.17	19.24
55-63	15.67	12.04

Table 6.2: The number of sessions per age group.

significant ($\chi^2 = 16.81$, $p = .002$). Furthermore, the Conover-Iman post-hoc test with Bonferroni correction showed that "Fingerprint" users ($M = 53.64$, $SD = 35.48$) had significantly more sessions than those using PIN ($M = 23.90$, $SD = 21.15$) and Swipe/None ($M = 30.88$, $SD = 24.01$). The distribution of the number of sessions across different unlocking mechanisms is shown in Figure 6.5.

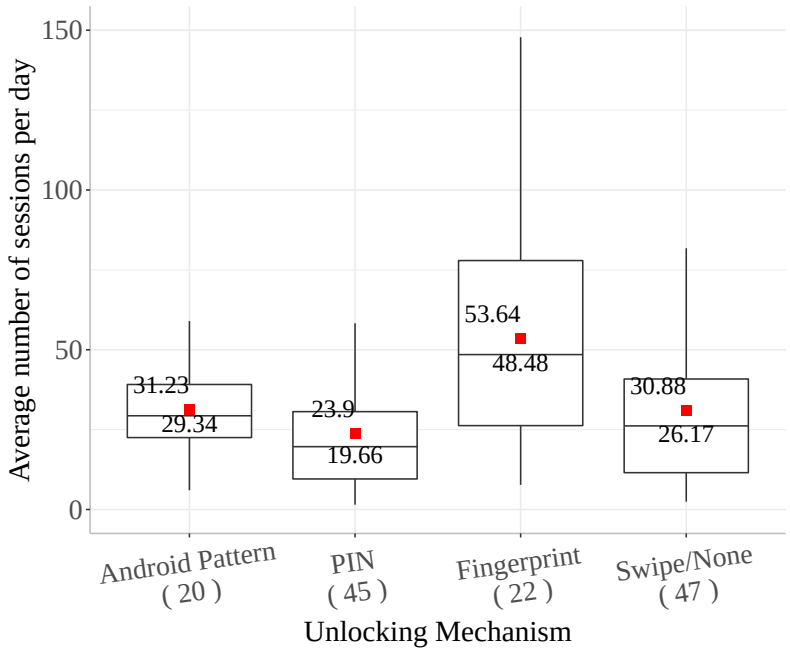

Figure 6.5: Average number of sessions per unlocking mechanism, $N = 134$ *(the three password participants were removed from the figure, as the sample size was too small)*. The red squares represent the means.

6.1.3 Daily Usage Lengths

In general, we found that younger participants tended to use their smartphones more frequently than older ones, i.e., in session counts. At the same time, an analysis of session length did not reveal any statistically significant difference between the age groups.

Figure 6.6 shows the daily averages for the total amount of time the participants used their smartphones per age group, and statistical analysis revealed that there was a significant difference (Kruskal-Wallis $\chi^2 = 15.58$, $p = .004$). A post-hoc Conover-Iman test (Bonferroni-corrected) revealed that participants from the age group "55-63" ($M = 106.89$, $SD = 89.48$) used their smartphones significantly less frequency of usage per day than participants from the age groups "19-24" ($M = 215.93$, $SD = 115.24$), "25-34" ($M = 251.24$, $SD = 161.70$), and "45-54" ($M =$

254.10, SD = 200.65).

Figure 6.7 shows the average amount of daily device usage time for participants using different unlocking mechanisms. Statistical analysis did not reveal any significant differences in total device usage time per day across various unlocking mechanism groups (Kruskal-Wallis $\chi^2 = 6.10$, $p = .19$).

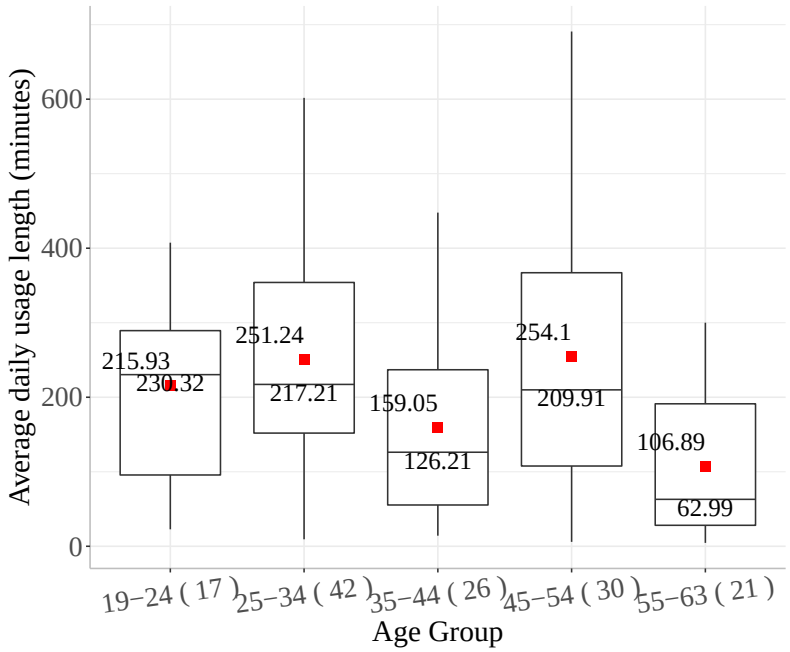

Figure 6.6: Average daily usage length per age group (minutes), $N = 136$. The red squares represent the means.

6.1.4 Error Rates

To assess how often participants made mistakes while unlocking, we calculated their error rates, i.e. how often their unlocking attempts failed. Since participants using the "Swipe/None" method cannot make errors while unlocking, we removed these participants from our analysis on error rates, which reduced the analysis sample to 90 participants. We tested both the impact of the type of unlocking

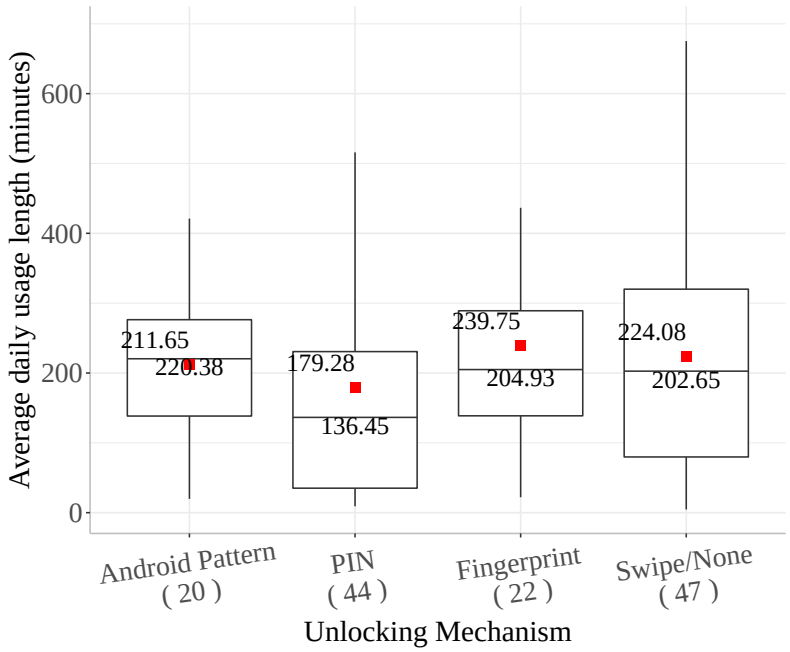

Figure 6.7: Average daily usage length per unlocking mechanism (minutes), $N = 133$ *(the three password participants were removed from the figure, as the sample size was too small)*. The red squares represent the means.

mechanism and age against the error rate. Statistical analysis did not reveal any significant differences among either age groups (Kruskal-Wallis $\chi^2 = 4.02$, $p = .40$) or unlocking mechanisms (Kruskal-Wallis $\chi^2 = 0.90$, $p = .82$). Figure 6.8 shows the distribution of error rates for participants based on their age groups and the unlocking methods used.

6.1.5 Auto/Manual Locks

We removed "Swipe/None" participants from our analysis of auto/manual locks, because in Android OS the autolock setting is not enabled for them. Furthermore, we removed 26 participants who disabled the autolock functionality manually, either at the beginning of the study or during the study. Overall, we analyzed varieties of locking behaviour based on the data collected from 64 participants.

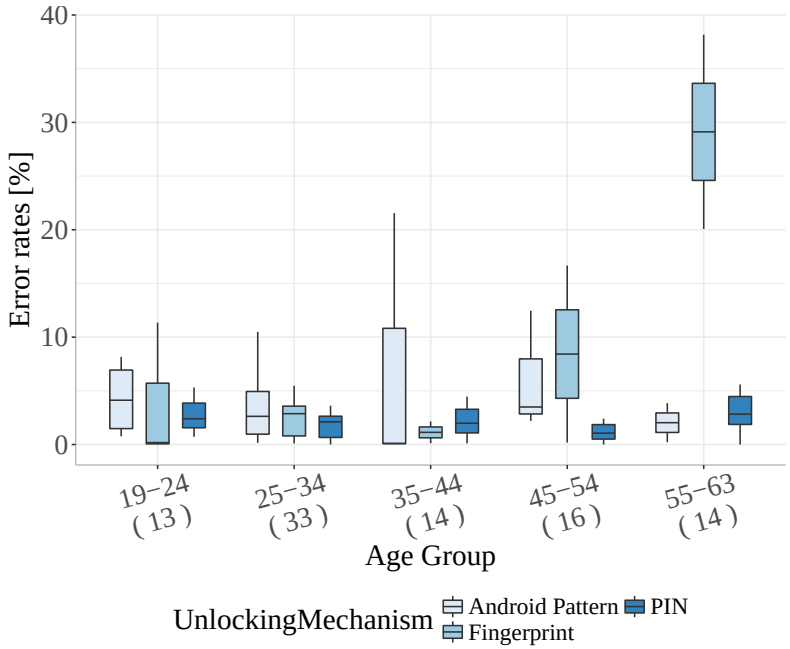

Figure 6.8: Error rates distribution among age groups, $N = 87$ *(the three password participants were removed from the figure, as the sample size was too small)*. Numbers below each age group represent the total number of participants from that group.

By default, the value of the autolock timeout is set to 5 seconds. Our analysis revealed that out of the 64 participants, 27 retained the default autolock timeout, while 18 reduced the autolock time to 0, which locks the device immediately after it enters sleep mode. Since we were not able to differentiate autolocks from manual locks when the timeout was set to 0, we further excluded those 18 participants from the analysis. This reduced our participant pool for this analysis to 46.

To evaluate whether the participants relied more on autolocks or manual locks, we calculated the percentages of autolocks over the total number of locks per day per user. Statistical analysis revealed that people's locking behaviour on manual lock vs autolock differed among the age groups (Kruskal-Wallis $\chi^2 = 11.74$, $p = .02$), but not across the unlocking mechanisms (Kruskal-Wallis $\chi^2 = 2.26$, p

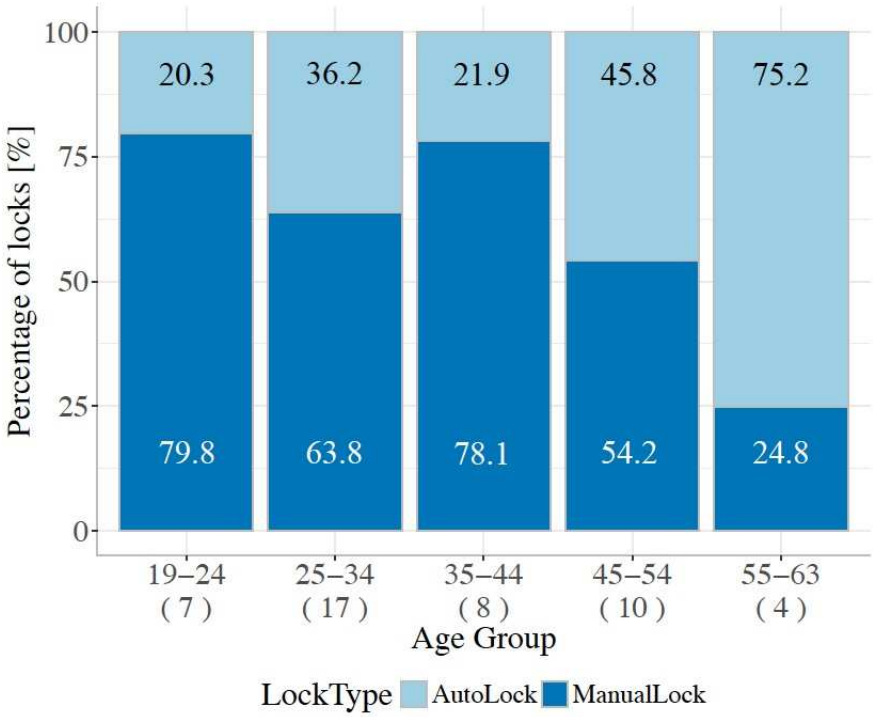

Figure 6.9: Distribution of locking types among age groups, $N = 46$. Numbers below each age group represent the total number of participants from that group.

= .52). Further analysis showed that participants in the age group "55-63" (M = 70.22, SD = 26.52) had significantly more autolocks than those in the younger age groups, "19-24" (M = 26.86, SD = 29.71), "25-34" (M = 32.54, SD = 27.22) and "35-44" (M = 21.40, SD = 14.44). The overall percentage of autolocks in each age group and unlocking mechanism group are presented in Figures 6.9 and 6.10, respectively.

6.2 Activity While Unlocking

Of the 137 participants, 119 chose to provide us with their activity data. To identify user activities during unlocks, we mapped their timestamps. We used one minute as the threshold, meaning that if no activity records occurred one minute before or

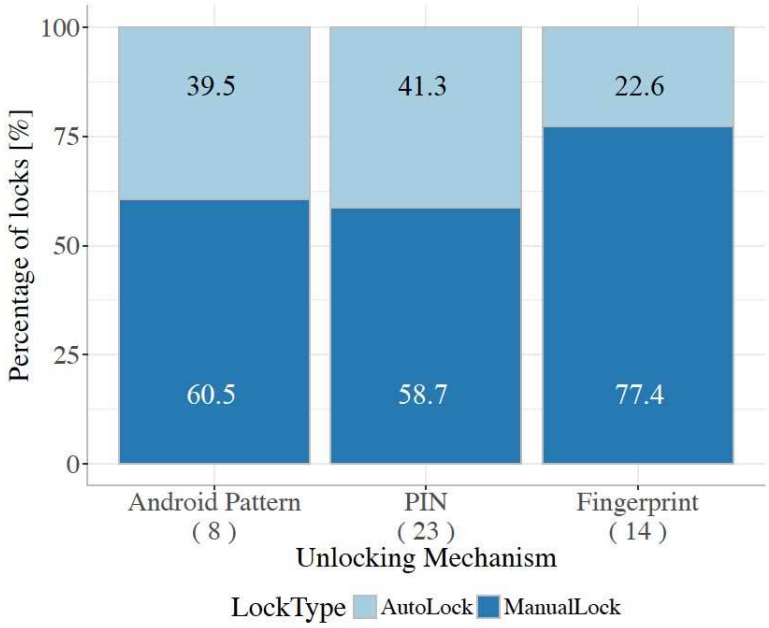

Figure 6.10: Distribution of locking types among unlocking mechanisms, $N = 45$ *(the only one password participant in this analysis was removed from the figure, as the sample size was too small).* Numbers below each unlocking mechanism represent the total number of participants from that group.

after the time when the unlock took place, then we considered that we had failed[1] to detect the type of activity associated with this unlock. When multiple activity records occurred within a one-minute timeframe, we selected the closest activity (time-wise) as the one that user had undertaken during the unlock.

Among all unlocking events collected for the 119 participants, we removed those that were tagged with "unknown"[2] activity and those for which we had failed to detect activities, which amounted to around 50.2% of the unlocks. Figure 6.11 shows the distribution of activity types over the remaining 49.8% of unlocks. We

[1]The potential reason for this could be that the sensor used by the API (Google's Activity Recognition API) for detecting the activity type was not started successfully while the unlock took place.

[2]Because there is not enough data for Google's Activity Recognition API to determine with significant confidence the activity the user is currently performing.

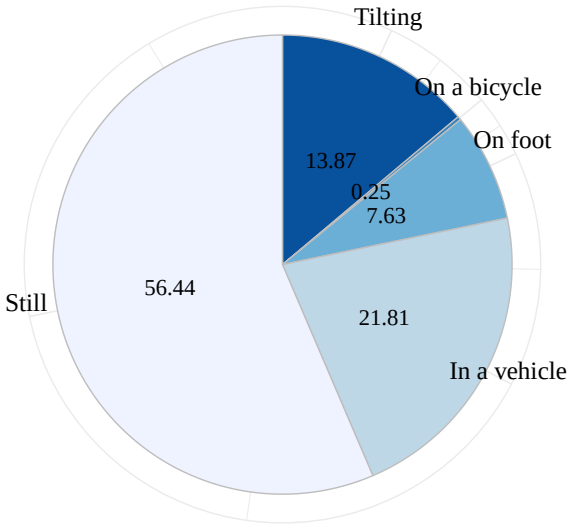

Figure 6.11: Distribution of activities during unlock, $N = 119$.

found that more than half of these unlocks took place when the participants' devices were still, whereas the other 44.55% of unlocks occurred while the device was moving. Only 0.25% of the unlocks happened while the participants were on a bicycle. The distributions of activities among the unlocking mechanisms did not differ notably from one another.

For the following, we removed the activity type "tilting", which we considered as not significant enough of a move to influence unlocking performance. Afterwards, we categorized all activities into two states, *still* and *move* (all activities other than *still*). As shown in Table 6.1, we found that age was a good predictor for whether authentication was used while the device was still or on the move, and the used authentication mechanism. Such findings are important as they help designers and developers to estimate the adoption rates of newly-designed authentication systems for users from different age groups and under different usage scenarios (*still* vs. *move*).

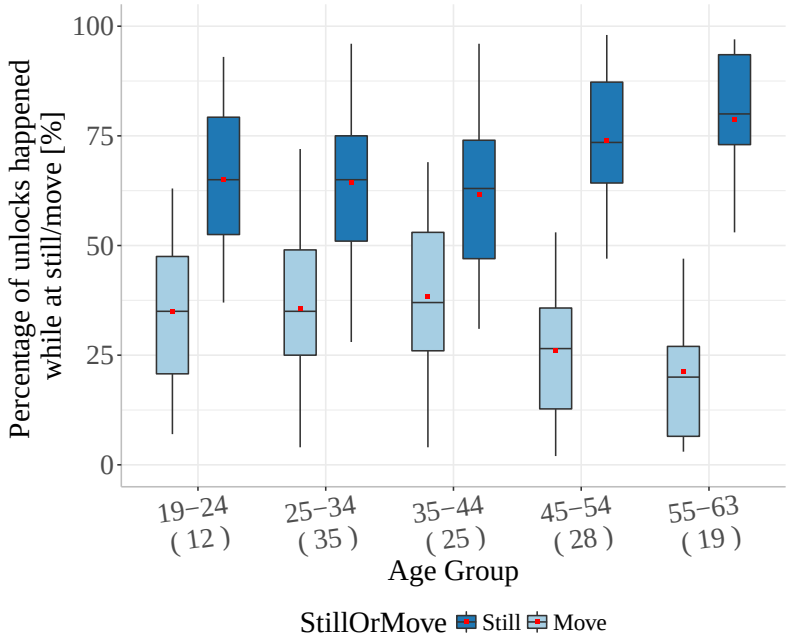

Figure 6.12: Distribution of unlocks that happened while at still/move sorted by age groups, $N = 119$. The red squares represent the means.

6.2.1 Error Rates While At Still vs On the Move

To assess how error rates correlate with activity types, we additionally excluded participants who used "Swipe/None" as errors are not possible with this mechanism.

We then counted the number of failed and total unlocking attempts that each participant performed while they were at "still" and "move" states. We then calculated the likelihood of making an error for each participant for each activity state. However, a paired-samples t-test did not reveal any significant differences in the likelihoods that participants would make an error while their devices were *still* ($M = 4.34$, $SD = 7.04$) or *moving* ($M = 4.82$, $SD = 7.26$); $t(73) = 1.27$, $p = .21$.

Variable	*Estimate*	*Std. Err.*	*z value*	*Odds Ratio*
Intercept	.428	.019	22.67*	1.53
Age=25-34	.084	.022	3.83*	1.09
Age=35-44	-.003	.025	-0.11	1.00
Age=45-54	.665	.025	26.96*	1.94
Age=55-63	.525	.031	16.74*	1.69

Table 6.3: Binary logistic regression model: whether authentication is used while at still vs on the move by age. An * denotes significance ($p < .05$).

6.2.2 How Age Predicts Whether Authentication is Used While At Still vs On the Move

To evaluate how often participants unlocked their smartphones while they were at still and moving states, we calculated the percentage of unlocks for each mobility type (*still* and *move*) per user over the entire study period. Figure 6.12 presents the detailed distribution of unlocks among the age groups and activity types. To further understand how age interacts with activity types (*still/move*), we fitted a binary logistic regression model to predict whether authentication will be used when users are at still and on the move based on age. We used age group "19-24" as the reference category for the model. The analysis revealed that age had a significant effect on predicting whether authentication happened while at still or on the move. Table 6.3 gives an overview of the fitted model.

6.2.3 How Age and Still/Move Predict Lock Types Used For Authentication

Another aspect worth looking at is the correlation between the lock type used for unlocks and other factors including age and activity states (*still*, *move*). We first excluded the lock type "Password" from the analysis, since very few (4) participants had used it. Afterwards, we fitted a multinomial logistic regression model to predict the lock type based on the participants' age and underwent

Variable	*Estimate*	*Std. Err.*	*z value*	*Odds Ratio*	*Estimate*	*Std. Err.*	*z*	*Odds Ratio*	*Estimate*	*Std. Err.*	*z*	*Odds Ratio*
	Swipe/None vs. Android Pattern				**Swipe/None vs. Fingerprint**				**Swipe/None vs. PIN**			
Intercept	-.140	.042	-3.30*	.87	.180	.039	4.60*	1.20	-.197	.043	-4.56*	.82
Age=25-34	-.438	.053	-8.21*	.65	.365	.046	7.89*	1.44	.698	.050	14.06*	2.01
Age=35-44	-1.341	.061	-22.09*	.26	-.859	.051	-16.93*	.42	-.274	.053	-5.22*	.76
Age=45-54	-2.195	.072	-30.69*	.11	-2.432	.068	-35.82*	.09	-.907	.055	-16.47*	.40
Age=55-63	-.684	.087	-7.87*	.50	-1.356	.095	-14.31*	.26	.827	.067	12.26*	2.29
Activity=Still	.214	.057	3.72*	1.24	.536	.052	10.32*	1.71	.547	.056	9.68*	1.73
Age=25-34:Activity=Still	-.020	.071	-0.29	.98	-.389	.061	-6.41*	.68	-.443	.065	-6.84*	.64
Age=35-44:Activity=Still	.493	.077	6.41*	1.64	-1.077	.069	-15.65*	.34	-.590	.069	-8.61*	.55
Age=45-54:Activity=Still	-.358	.088	-4.07*	.70	-1.262	.087	-14.48*	.28	-1.005	.070	-14.40*	.37
Age=55-63:Activity=Still	.451	.102	4.42*	1.57	-.283	.111	-2.55*	.75	-1.324	.085	-15.60*	.27

Table 6.4: Multinomial logistic regression model: how age and still/move states predict lock type used for authentication. An * denotes significance ($p < .05$).

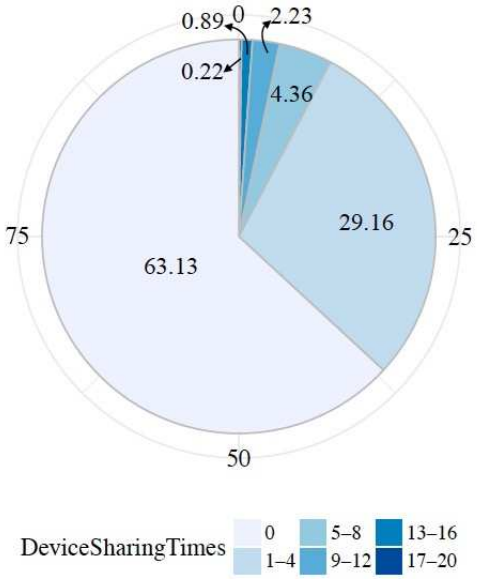

Figure 6.13: Distribution of answers to: "Since last week, how many times have you shared this device with others?"

activities, with interaction between age and activity types. We used lock type "Swipe/None", the age group "19-24" and the activity state "move" as the reference categories for the model. In general, our results show that all the main effects were significant, in addition to the interaction effect between age and activity states (*still/move*). Table 6.4 provides an overview of the fitted model.

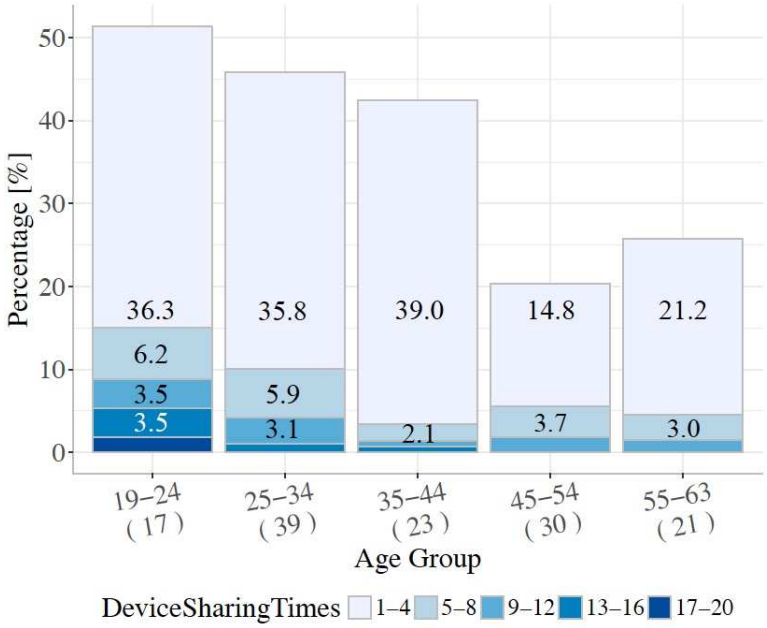

Figure 6.14: Distribution of answers to: "Since last week, how many times have you shared this device with others?" among age groups, with "0" excluded from the figure, $N = 130$. The percentages shown in y-axis are the percents of received responses of device sharing times for all participants among each age group.

6.3 Device Sharing

Our study app presented participants with the device sharing survey once a week. Out of our 137 participants, 130 answered at least one of the device sharing surveys. In total, we received 897 responses. On average, each participant answered 7 such questionnaires (SD = 2.4, Min = 1, Max = 17). Overall, out of the 130 participants, 81 reported that they regularly shared their device with others. In all the device sharing survey responses that we received, participants reported over 60% of the time that they had not shared their devices with others during the preceding week. The detailed distribution of device sharing times is presented in Figure 6.13.

Figure 6.14 shows how responses on device sharing times (greater than 0) are distributed among the age groups. To evaluate how preferences for sharing

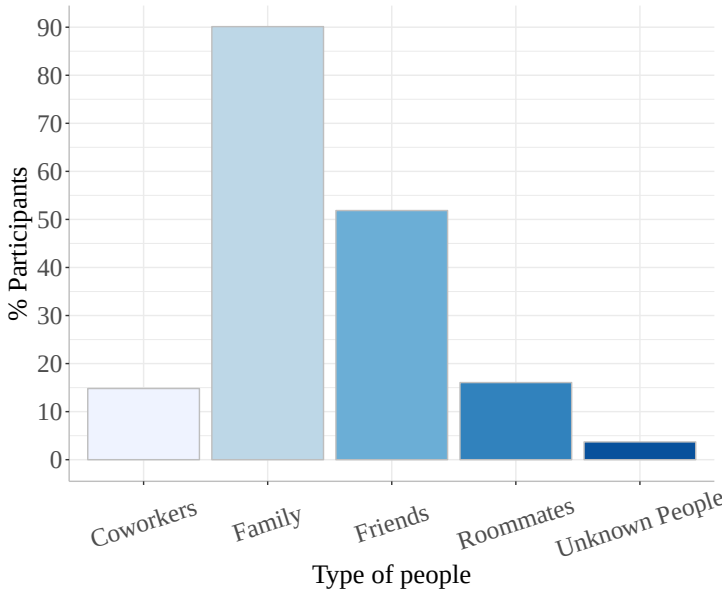

Figure 6.15: Distribution of device sharing with various groups of people, $N = 81$

differ with age, we further divided all responses into two categories: *not shared* (responses on sharing for 0 time) and *shared* (all responses except *not shared*). We then calculated the corresponding percentages of all *shared* responses for each participant. Interestingly, while a Kruskal-Wallis test revealed that participants differed significantly in their device sharing preference among the age groups (χ^2 = 11.73, p = .02), further analysis (Bonferroni-corrected Conover-Iman post hoc) did not reveal any significant difference in the percentages between any two age groups.

As shown in Figure 6.15, family and friends are the top two types of people with whom participants regularly shared their devices, followed by roommates and co-workers. Interestingly, about 3.7% of the 81 participants also reported that they shared their devices with unknown people.

Overall		19-24		25-34		35-44		45-54		55-63	
App name	*rating*	*App name*	*rating*	*App name*	*rating*	*App name*	*rating*	*App name*	*rating*	*App name*	*rating*
Chrome	1.74	Chrome	2.19	Chrome	1.82	Chrome	1.82	Facebook	1.38	Chrome	1.86
Facebook	1.87	Gmail	2.73	Gmail	2.29	Contacts	1.73	Chrome	1.18	Contacts	2.77
Google Play	1.51	Snapchat	3.38	Facebook	1.90	Facebook	1.99	Messenger	1.50	Settings	1.95
Settings	1.53	YouTube	2.00	Google Play	1.71	Settings	1.39	Google Play	1.16	Facebook	2.46
Gmail	2.24	Settings	1.97	Contacts	2.07	Google Play	1.53	Contacts	1.32	Google Play	1.58

Table 6.5: Top 5 most-used apps and the corresponding averaged sensitivity ratings. Sensitivity level has been collected using a 5-point Likert scale (1: Not at all, 2: Slightly, 3: Moderately, 4: Very, 5: Extremely.) with question: "Please rate for each app on the list: If this person were watching your screen, and you were using this app right now, how much would it affect your privacy?"

6.4 App Usages and Sensitivities

In this section, we report on the apps that participants accessed the most during the study, with sensitivity ratings for these apps. We also present results on how participants from different age groups perceived the sensitivities of their apps.

We excluded the data of the 8 participants who had not granted permission for our study app to collect their application details. In addition, we excluded special system apps, such as *launcher* and *systemui*, from the analysis, since these apps form a part of Android OS, which the user did not actively open or access.

In total, we received 18,496 app sensitivity survey responses. On average, each participant answered 136 such questionnaires ($SD = 102.9$, $Min = 1$, $Max = 732$). Our analysis showed that our participants used 2,976 unique apps overall. On average, each participant used 72 unique apps during our study, ranging from 10 to 200 ($SD = 36$, $Median = 67$).

We then identified the most-used apps. To do this, we first aggregated the number of days on which each application was launched by each participant. Afterwards, we summed up the days used for each app by all the participants; we then considered the sum as the total number of days that this app was used during the study, and sorted all the apps by this number. Table 6.5 shows the top 5 most-used apps for all the participants overall and across each age group.

To evaluate how sensitive the apps are from the participants' points of view,

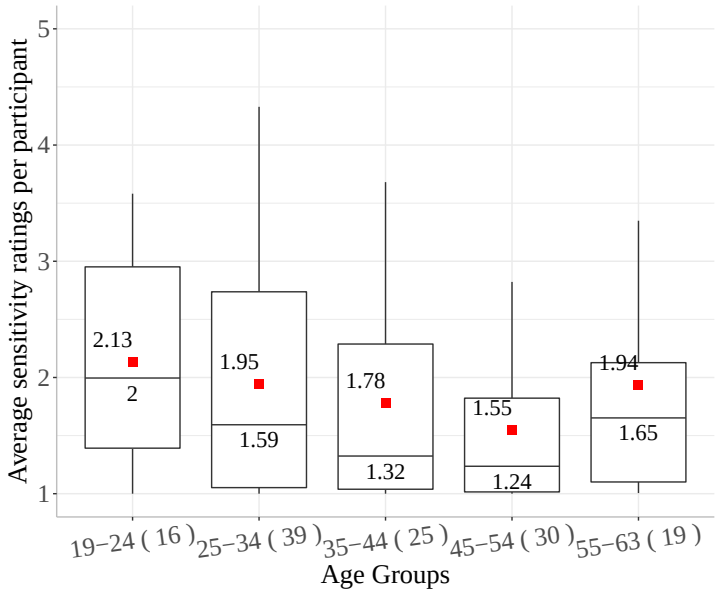

Figure 6.16: Participants' perceptions of app sensitivity (5-point Likert-type scale (1: Not at all, 2: Slightly, 3: Moderately, 4: Very, 5: Extremely.)) across age groups, $N = 129$. The red squares represent the means.

we calculated the average sensitivity ratings for each application across all the participants, who rated the apps in their specific groups. The sensitivity ratings for the top 5 most-used apps in each group are also shown in Table 6.5.

Overall, Chrome was the most used app among all the age groups, while participants from the "19-24" group rated it with the highest sensitivity rating, compared with the other groups. The top 5 most-used apps for the participants from the "35-44" and "55-63" groups were identical with the exception of order.

To assess how participants' perceptions of app sensitivity differ with age, we calculated the average sensitivity ratings among all the rated applications for each participant. On average, the participants rated 1.83 for all apps (SD = 0.94, Min = 1, $Median$ = 1.43, Max = 4.96). Figure 6.16 shows the distribution of the average sensitivity ratings per participant across the age groups. However, a Kruskal-Wallis test did not reveal any significant differences among the age groups (χ^2 = 7.04, p

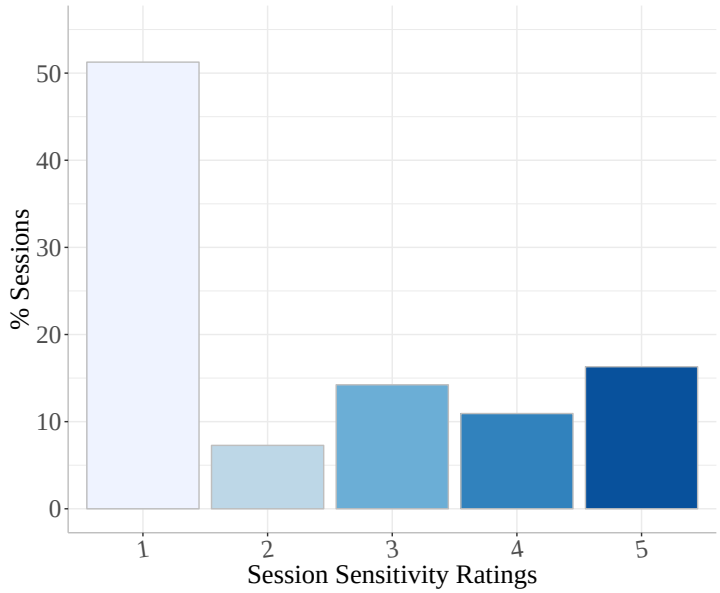

Figure 6.17: Distribution of session sensitivity ratings (5-point Likert-type scale (1: Not at all, 2: Slightly, 3: Moderately, 4: Very, 5: Extremely.)), $N = 129$.

= .13).

We also checked how many sessions included apps that the respective user rated highly-sensitive (4 or 5 on a 5-point Likert-type scale; 1: Not at all, 2: Slightly, 3: Moderately, 4: Very, 5: Extremely). To do this, we defined and calculated the sensitivity ratings for each session as the maximum of all app ratings collected from it. As shown in Figure 6.17, only 27.2% of the 16, 893 rated sessions were considered highly-sensitive.

Chapter 7

Discussion

The results presented in this paper provide initial insights into how individuals' smartphone authentication patterns differ by age group, thus, extending on previous studies [7, 8, 12, 13].

7.1 Age Makes a Difference

First of all, our data shows that smartphone usage patterns indeed differ based not only on unlocking mechanisms, but also on age. Our results quantitatively provide support for previous claims that demographic differences such as age should be taken into account while designing new authentication systems for smartphones [8].

For instance, we found that participants in older age groups interacted with their devices significantly less frequently than younger groups. They also interacted with their devices for a significantly shorter amount of time each day. This means two things: 1) they are exposed to the authentication mechanism less frequently than younger groups and 2) authentication takes up a greater portion of the overall interaction with their devices.

In addition, age also significantly correlated with the locking behaviour of our participants. We found that older groups relied more on autolocks than the other groups. Overall, the "55-63" participants were more than three times as likely to

use the autolock feature than participants from the “19-24” and “35-44” groups, and about two times more likely than those from the “25-34” and “45-54” groups. As the autolock feature presents a trade-off between security and usability, future research into the popularity of the autolock function among older users and their awareness of the respective security implications is necessary.

Another interesting result to highlight here is that age was a good predictor for whether authentication was used with a moving device vs a still device. In combination with those two states, age was also a good predictor for which authentication system is used. This result further helps predict and evaluate whether certain authentication systems would likely be adopted among various age groups, which is important for designing future systems.

7.2 No One-Fits-All Solution

We found that participants using different unlocking mechanisms use their devices differently. Specifically, Fingerprint users tend to have more device sessions than other groups, including the PIN and Swipe/None groups. In addition, PIN and Android unlock pattern users were found to be two times more likely to rely on autolock features than Fingerprint users.

Different usage across the authentication systems has been shown before [9]. Our results extend these previous findings by showing that these usage differences hold true for more diverse samples as well, including for older age groups. We are also the first to extend these findings to Fingerprint users.

Our numbers do not allow us to infer the reasons for why users pick their respective authentication system (or the direction of causality). It is possible that people select their authentication system with respect to their needs (e.g., shorter authentication time due to more sessions, which was more likely in younger age groups). Another possible interpretation could be that participants adapt their behaviour to the authentication system they have chosen.

That said, while we cannot make claims about the reasons for various user choices, the results of our study indicate that offering a selection of different

authentication systems for smartphones seems to be the right thing to do. Users have different preferences and this approach allows them to choose the system that best fits their needs and the way they are using their devices.

7.3 Design and Evaluate for Movement

Our results show that, overall, about 50% of unlocks happened when the user was moving (we detected the phone movement; hence, it was safe to assume that the participant was moving). This was a common observation across all age groups.

One might hypothesize that participants were more likely to make unlocking errors while moving. Our dataset, however, shows that this was unlikely. The numbers of failed unlocking attempts when participants' devices were still and when they were moving were very similar across the age groups. This might indicate that current unlocking mechanisms are robust against a moving state.

A consequence of this and the fact that people do regularly authenticate themselves while moving, means that newly-designed unlocking mechanisms should be as robust in the presence of movement as the existing mechanisms, in order to be acceptable for smartphone unlocking. This requirement not only influences how unlocking systems should be designed (e.g., constant eye contact might not be possible) but also how they are evaluated in studies (in-lab vs. in the wild).

A potential caveat to our findings is that our data collection instrument failed to detect the activity types for around 50% of authentication events, due to technical limitations (the activity recognition API reported "unknown" or failed to detect the undertaking activity). As such, readers should take this into consideration while interpreting these results.

7.4 Device Sharing is Common

A majority of the participants in our study (around 60%) had shared their devices with others. While most of them reported that they had shared their devices with family members and friends, there is a small group of people (3) who had shared

their devices with unknown people. Our Android results corroborate with the iPhone sharing patterns reported by Cherapau et al. [3], who found that 60% of iPhone users have also shared their passcodes with others (with partners, family, and friends being their top three sharing recipients).

Furthermore, while we found that age significantly correlates with the number of sharings taking place per week, it was still a very common task across all the age groups.

This means that authentication systems will benefit from being flexible in allowing sharing (to a certain extent) while continuing to remain protected. For example, a behavioural biometrics system that checks for anomalies to lock a device should be designed in a way that it does not make sharing impossible. An authentication mechanism that can be shared among a trusted set of people without revealing their own secrets could also be useful for people who tend to share their phones with others.

7.5 Most Sessions Have Low Sensitivity

Our results show that, in the majority of cases, participants accessed apps that they considered to have low sensitivity (roughly 72% of sessions). This is similar to what was found by Harbach et al. [9]. In their study, in most sessions, participants accessed data that they considered to have little to no sensitivity. Additionally, we found that, while participants' individual perception of app sensitivity varies (from 1 to 4.96), it does not correlate with age.

An important caveat here is that this data relies on self-reporting, and as shown in previous work, users often underestimate the actual privacy risks associated with smartphone access [5].

Nonetheless, this data provides further support for the claim that a significant portion of smartphone functionality could be made available without the need for authentication. This is especially important due to the large number of sessions (and time) that users spend authenticating themselves on their devices. Smartphone manufacturers have been moving along that direction recently (e.g., activating

alarms without logging in to the phone). We argue that identifying more of these use cases could help to further reduce the amount of time that smartphone users spend in unlocking their devices, without compromising their privacy and security.

Chapter 8

Limitations

While we have worked to mitigate any potential problems with the setup of this study, a few limitations must be kept in mind when processing the present results.

First, our server was silently down for 10 days at the end of May, when most of our participants were supposed to complete the study. This unexpected issue forced us to drop unrecoverable data for certain participants. This could have potentially biased our study sample. However, we do not consider this a major threat to the validity of our results, since we did not find statistically significant differences between the distributions of our sample demographics and the demographics of the US smartphone ownership reported by the Pew Research Center.

We measured session length as the time between screen on/off (with keyguard removed) events. Since, with this measurement, we could not remove the screen off timeout from the session length calculations, the reported session lengths could potentially have been longer than the actual ones. Our comparisons of session lengths among the groups might also be biased as a result, as different participants would have set the screen off timeout to different values, varying from 0 seconds to 24 hours.

Due to technical limitations, we were unable to programmatically detect the unlocking mechanisms and states of autolock settings employed by participants using Android 6 (64 users) and 7 devices (31 users). Therefore, we asked those

participants to report their unlocking mechanisms (Android 6 and above users) and the autolock settings (Android 7 and above) in the entry and exit surveys. As a consequence, the correctness of our collected data partially depends on user recollection. However, the results give us no reason to believe that they would have incorrectly reported this data. In addition, the two other mini-questionnaires (the device sharing survey and the contextual survey) that we presented during the study also relied on self-reported data and may have impacted on the correctness of our results, despite being as close to the actual events as possible. Furthermore, our device sharing questions did not define the concrete "sharing" scenarios, i.e., sharing content on the smartphone screen with others, handing over device under supervision, or handing over device unsupervised, and the interpretation of sharing depends on the participants.

On average, the confidence level of the detected activities is 73.37 (SD = 25.20, Min = 21, $Median$ = 75, Max = 100). Therefore, our activity analysis results are highly dependent on the accuracy of the activity recognition API.

Chapter 9

Conclusion

In this research, we provide the first detailed investigation of how age of smartphone users correlates with their unlocking behaviour. We conducted a longitudinal field study with 137 participants from North America, who installed our study app on their Android phones and ran it for at least 60 days. The results of our study suggest that age does correlate with certain patterns of smartphone use related to unlocking. In particular, we observed that older users interacted with their devices for significantly shorter amounts of time than younger ones, while, at the same time, they relied more on autolock features. Overall, all participants spent similar amounts of time interacting with their devices per user session, whereas younger groups were likely to use their devices for significantly more time in total each day. We highlight that, when designing new authentication systems, varying age-related usage patterns should be taken into consideration.

We also show that about 50% of unlocks happened when smartphone users were in a state of motion. This indicates that it is important for newly designed unlocking mechanisms to be robust against activities involving movement.

Additionally, we found that user interaction differs depending on the unlocking mechanisms they use. Fingerprint users were found to interact with their device significantly more frequently than PIN and Swipe/None users. Our results did not reveal significant differences in other usage data across all types of unlocking

mechanism.

We found that about 62% of participants shared their devices with others, with family members and friends being the top two device sharing targets. Finally, we identified the top five most-accessed apps for all participants and for participants from each age group. Our participants provided low sensitivity ratings for the most-used apps on their devices, which indicates that users are likely to underestimate the value of data stored on their devices [5]. Last but not least, our study did not show that participants' general perception of app sensitivity varies by age.

We see two promising areas for future work in this field. First, now that we have a better understanding of how users' (un)locking behaviours vary under different circumstances such as age, it would be interesting to find out why these differences exist. Furthermore, exploring the understanding of security and usability trade-offs between age groups and unlocking mechanisms can potentially shine further light on users' decision making processes.

Bibliography

[1] D. Buschek, F. Hartmann, E. von Zezschwitz, A. De Luca, and F. Alt. Snapapp: Reducing authentication overhead with a time-constrained fast unlock option. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 3736–3747, New York, NY, USA, 2016. ACM.

[2] P. R. Center. Smartphone ownership 2013, 2013.

[3] I. Cherapau, I. Muslukhov, N. Asanka, and K. Beznosov. On the impact of touch ID on iphone passcodes. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 257–276, Ottawa, 2015. USENIX Association.

[4] M. Cherubini and N. Oliver. A refined experience sampling method to capture mobile user experience. In *In Presented at the International Workshop of Mobile User Experience Research part of CHI*, pages 1–12, Boston, MA, USA, 2009. ACM.

[5] S. Egelman, S. Jain, R. S. Portnoff, K. Liao, S. Consolvo, and D. Wagner. Are you ready to lock? In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 750–761, New York, NY, USA, 2014. ACM.

[6] Google. Detected activity types, 2017.

[7] M. Harbach, A. De Luca, and S. Egelman. 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*, CHI '16, pages 4806–4817, New York, NY, USA, 2016. ACM.

[8] M. Harbach, A. De Luca, N. Malkin, and S. Egelman. Keep on lockin' in the free world: A multi-national comparison of smartphone locking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 4823–4827, New York, NY, USA, 2016. ACM.

[9] M. Harbach, E. von Zezschwitz, A. Fichtner, A. De Luca, and M. Smith. It's a hard lock life: A field study of smartphone (un)locking behavior and risk perception. In *Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 213–230, Menlo Park, CA, 2014. USENIX Association.

[10] E. Hayashi, S. Das, S. Amini, J. Hong, and I. Oakley. Casa: Context-aware scalable authentication. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, pages 3:1–3:10, New York, NY, USA, 2013. ACM.

[11] E. Hayashi, O. Riva, K. Strauss, A. J. B. Brush, and S. Schechter. Goldilocks and the two mobile devices: Going beyond all-or-nothing access to a device's applications. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS2012, pages 2:1–2:11, New York, NY, USA, 2012. ACM.

[12] D. Hintze, R. D. Findling, M. Muaaz, S. Scholz, and R. Mayrhofer. Diversity in locked and unlocked mobile device usage. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, UbiComp '14 Adjunct, pages 379–384, New York, NY, USA, 2014. ACM.

[13] A. Mahfouz, I. Muslukhov, and K. Beznosov. Android users in the wild: Their authentication and usage behavior. *Pervasive and Mobile Computing*, 32:50–61, 2016.

[14] D. Marques, I. Muslukhov, T. Guerreiro, L. Carriço, and K. Beznosov. Snooping on mobile phones: Prevalence and trends. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 159–174, Denver, CO, 2016. USENIX Association.

[15] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov. Know your enemy: The risk of unauthorized access in smartphones by insiders. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services*, MobileHCI '13, pages 271–280, New York, NY, USA, 2013. ACM.

[16] P. Northcraft. Android: The Most Popular OS in the World, 2014. Last Accessed: June 2017.

[17] O. Riva, C. Qin, K. Strauss, and D. Lymberopoulos. Progressive authentication: Deciding when to authenticate on mobile phones. In *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*, pages 301–316, Bellevue, WA, 2012. USENIX.

[18] D. Van Bruggen, S. Liu, M. Kajzer, A. Striegel, C. R. Crowell, and J. D'Arcy. Modifying smartphone user locking behavior. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, pages 10:1–10:14, New York, NY, USA, 2013. ACM.

Appendix A

Survey Questions

A.1 Device Sharing Survey

1. Would you be concerned if someone in your social circle is able to access this device **without your permission**?
   (a) YES
   (b) NO
2. Since last week, HOW MANY TIMES have you shared this device with others? (provided with a breakdown number list that user can specify themselves)
3. What kind of people did you share this device with? (Choose all that apply.)
   (a) Friends
   (b) Roommates
   (c) Family
   (d) Co-workers
   (e) Unknown people
   (f) Other (please specify)

A.2 Contextual Survey

1. Who can see the content of your screen right now?
   (a) Unknown person
   (b) Known person
   (c) Both
   (d) Nobody
2. Please rate how likely it is that someone is watching your screen right now.
   (a) Very unlikely
   (b) Unlikely
   (c) Neutral
   (d) Likely
   (e) Very likely
3. Please rate for each app on the list: If this person were watching your screen, and you were using this app right now, how much would it affect your privacy?
   (a) Extremely
   (b) Very
   (c) Moderately
   (d) Slightly
   (e) Not at all

Appendix B

Data Collection Visualization

Figure B.1 visualizes the data collection process of this research project from December 8, 2016 to August 10, 2017. The x-axis represents the date that the data was collected, and the y-axis shows the numeric ID of each participant. Each dot in the figure represents the data collected for a participant on a specific day. Each horizontal line represents the entire data collection process for each participant. As shown in Figure B.1, all participants are categorized to four groups, namely "Completed", "Failed to recover", "Partial data", "Withdrew". While "Completed" means the participants successfully provided us with data for 60 (or more) days, "Withdrew" means the participants stopped participating before the end of the study. "Failed to recover" participants were those whose data were lost due to the back-end server issue (our server was silently down for 10 days) and were not recovered by the end of the study (August 10, 2017). Participants who were categorized to the "Partial data" group were those whose data were not sent successfully to the server due to unknown issues related to their devices.

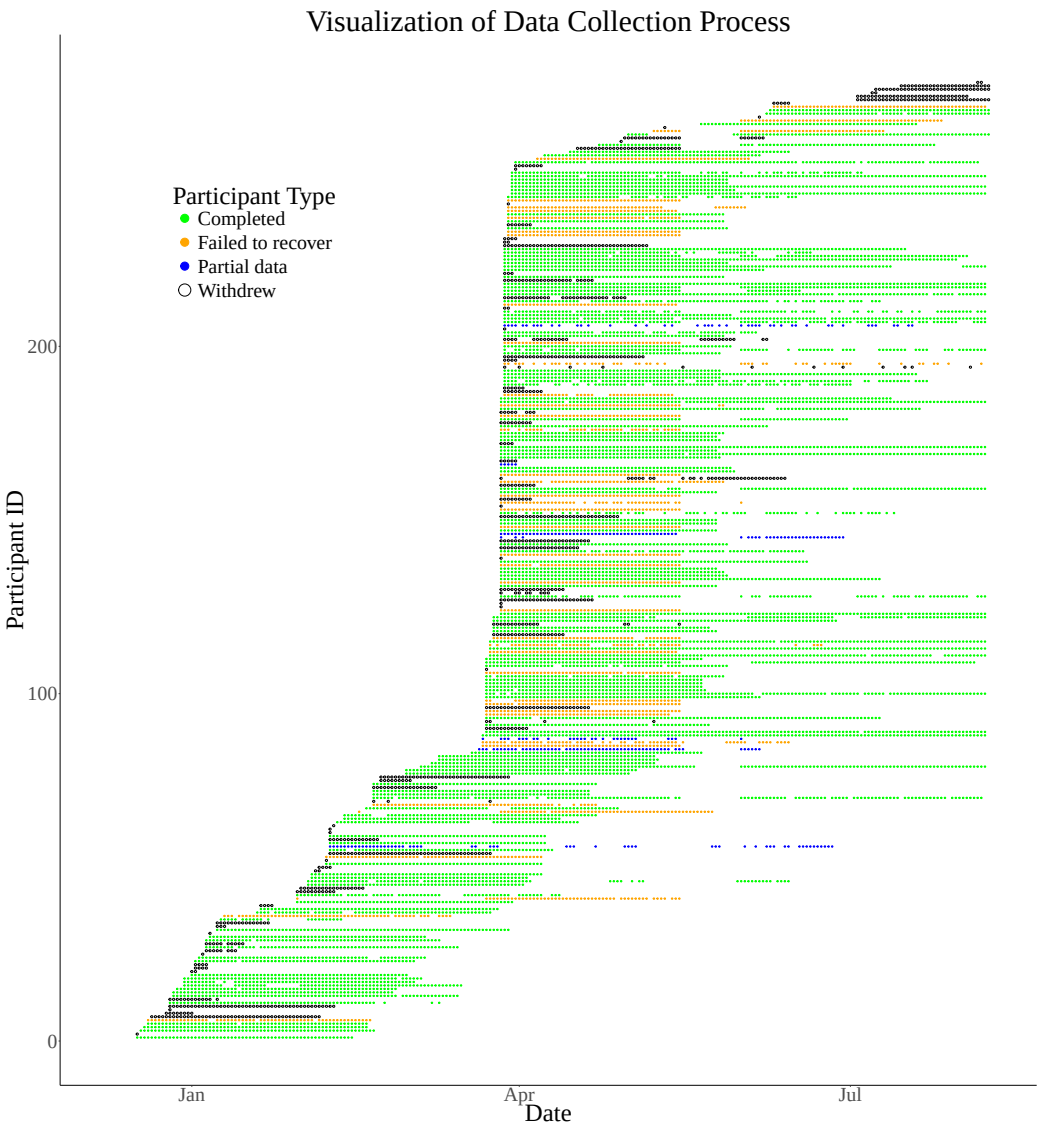

Figure B.1: A Visualization of Data Collection Process from December 8, 2016 to August 10, 2017, $N = 276$. Total # of Participants Who Completed the Study: 142; Total # of Participants from Whom We Failed to Recover Data: 43; Total # of Participants from Whom We Received Partial Data: 7; Total # of Participants Who Withdrew from the Study: 84