Can Machine Learning Help People Configure Their Mobile App Privacy Settings?

Bin Liu

CMU-ISR-19-105
December 2019

Institute for Software Research
School of Computer Science
Carnegie Mellon University
Pittsburgh PA 15213

Thesis Committee:
Norman Sadeh (Chair)
Alessandro Acquisti (Heinz College)
Lorrie Cranor
Florian Schaub (University of Michigan)
Nina Taft (Google Inc.)

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Societal Computing.

This research was sponsored by the National Science Foundation under grants CNS-0905562, CNS-1012763, CNS-1330596, and SBE-1513957, as well as by DARPA and the Air Force Research Laboratory, under agreement number FA8750-15-2-0277. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. Additional funding has also been provided by Google through a Google Faculty Research Award and the Google Web of Things Expedition and in part through a grant from the CMU-Yahoo! InMind project, as well as by the Carlsberg Foundation. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the Air Force Research Laboratory, the National Science Foundation, the U.S. Government, Google or Yahoo!.

Keywords: Mobile Privacy, Mobile Apps, Android Permissions, Assistants

Abstract

Technologies such as mobile apps, web browsers, social networking sites, and IoT devices provide sophisticated services to users. At the same time, they are also increasingly collecting privacy-sensitive data about them. In some domains, such as mobile apps, this trend has resulted in an increase in the breadth of privacy settings made available to users. These settings are necessary because not all users feel comfortable having their data collected by some of these technologies. On mobile phones alone, the sheer number of apps users download is staggering. The variety of sensitive data and functionality requested by these apps has led to a demand for much more specific privacy settings. The same is true in other domains as well, such as social networks, browsers, and various IoT technologies. The result of this situation is that users feel overwhelmed by all of the settings available to them, and are thus unable to take advantage of them effectively.

This dissertation examines whether machine learning techniques can be utilized to help users manage an increasingly large number of privacy settings. It specifically focuses on mobile app permissions. The research presented herein aims to simplify people's tasks in regard to managing their large number of app privacy settings. We present the methods we used for developing models of users' privacy preferences, and describe the interactive assistant we designed based on these models to help users configure their settings using personalized recommendations. The objective of this work is to alleviate the burden placed on users while increasing alignment between a their preferences and the privacy settings on their phones.

This dissertation details three different studies. Specifically, in the first study, we used a dataset of mobile app permission settings obtained from over 200K Android users, explored different machine learning models, and analyzed different combinations of features to predict users' mobile app permission settings. The study includes the development and evaluation of profile-based models as well as individual prediction models. It also includes simulation studies, wherein we explored the viability of different interactive configuration scenarios by testing different ways of combining dialogue inputs from users with recommendations based on machine learning models. The results of these simulations suggest that by selectively prompting users to indicate how they would like to configure a relatively small percentage of their permission settings, it is possible to accurately predict many of their remaining permission settings. Another significant finding of this first study is that a relatively

small number of privacy profiles derived from clusters of like-minded users can help predict many of the permission settings that users in a given cluster prefer.

The second study was designed to validate these findings in a field study with actual users. We designed an enhanced version of Android's permission manager and collected rich information on users' actual app permission settings. While results from this study involve a much smaller number of users, they were obtained using privacy nudges designed to increase user awareness of data being collected about them and as a result also their engagement with their permission settings. Using data collected as part of this study, we were able to generate and analyze privacy profiles built for groups of like-minded users who exhibited similar privacy preferences. Results of this study confirm that a relatively small number of profiles (or clusters of users) can capture s large percentahe of users' diverse privacy preferences and help predict many of their desired privacy settings. They also indicate that privacy nudges can be very effective in motivating users to engage with their permission settings and in deriving privacy profiles with strong predictive power.

In the third study, we evaluated our profile-based preference models by developing a privacy assistant that helps users configure their app permission settings based on the developed profiles from our second study. We report on the results of a pilot study (N=72) conducted with actual Android users who used our privacy assistant on their smartphones while performing their regular daily activities. The results indicate that participants accepted 78.7% of the recommendations made by the privacy assistant and kept 94.9% of these settings on their phones over the following six days, all while receiving daily nudges designed to motivate them to further review their settings. The dissertation also discusses the privacy profiles designed for this research and identifies essential attributes that separate people associated with different profiles (or clusters). A refined version of the Personalized Privacy Assistant was released to the Google Play store and used to collect some additional data.

In summary, through a series of three studies, this dissertation shows that using a small number of privacy decisions made by a given smartphone user, it is often possible to predict a large fraction of the mobile app permission settings this user would want to have. The dissertation further shows how we have been able to effectively operationalize this finding in the form of personalized privacy assistants that can help users configure mobile app permission settings on their smartphones.

Acknowledgments

I would like to thank everyone who supported me, mentored me, and helped me.

I would like to thank my advisor, Prof. Norman Sadeh, who patiently provided me detailed guidance on research training, design, and thinking. I am grateful for his valuable support in this research and in my exploratory work at Carnegie Mellon.

I would like to thank my thesis committee members, Prof. Alessandro Acquisti, Prof. Lorrie Cranor, Prof. Florian Schaub, and Dr. Nina Taft, who closely collaborated with our research team. I have learned a great deal from this joyful experience.

I would like to give thanks to my parents, my wife, and all of my other loved ones from whom I received steady love and emotional support throughout my academic career.

I would like to also thank my labmates and colleagues Dr. Jialiu Lin, Prof. Hazim Almuhimedi, Prof. Yuvraj Agarwal, Prof. Sebastian Zimmeck, Prof. Shomir Wilson, Prof. Lujo Bauer, Prof. Anupam Das, Dr. Mads Schaarup Andersen, Dr. Martin Degeling, Dr. Bin Liu, Dr. Yuanyuan Feng, Dr. Gaurav Misra, Dr. Piotr Mardziel, Joshua Gluck, Aerin Zhang, Peter Story, Daniel Smullen, Linda Moreci, Connie Herold, and all other people who shared their research expertise, excellent ideas, kindness, and support so that I could enjoy a friendly and collaborative environment doing research in this exciting domain.

Contents

Chapter 1

Introduction

1.1 Information Privacy

While different definitions have been proposed[91], information privacy is by and large about the ability of people to control what information about them is being collected and how it is used.

As information services become more sophisticated, they also increasingly rely on the collection and use of personal information. For instance, many smartphone apps collect their users' location, whether to tailor some of the content they show users or to help advertisers more effectively target these users. Internet of Things devices such as Amazon Echo or Nest further add to this trend and enable technology providers to access ever richer streams of personal data. This collection of sensitive data creates a complex collision of interests between different stakeholders such as users, service providers, and other third parties that are given access to this data.

Information privacy has been addressed in many different legal documents such as the Privacy Act of 1974[83], HIPAA[2], COPPA[3], and CalOPPA[4]. As regulatory guidelines for protecting the privacy of people and their data, the *Fair Information Practices Principles* (FIPPs) started generating attention in 1973[90]. The concept of FIPPs has been evolved and mirrored in many laws or policies internationally. For example, the Organization for Economic Cooperation and Development (OECD) proposed eight principles in 1980[1] including collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability. US Federal Trade Commission (FTC) identified five core principles of privacy protection in 1998[5], namely notice/awareness, choice/consent, access/participation, integrity/security, and enforcement/redress. Later this list was elaborated into eight principles[42]. As the internet and other digital services evolve, we may see more refined versions or implementations of fair information practices. For example, the European Union adopted the "General Data Protection Regulation" (GDPR)[6], which was adopted in 2016 and was put in effect in 2018. Under GDPR, parties who collect personal data have to provide intelligible and easy-

to-access consent to users, and obtain users' consent. Users also have control over keeping or erasing their personal data that has been collected.

In this dissertation, we primarily focus on the commonly referenced notion of "choice." Consumers should be given sufficient awareness of data collection as well as options to control who has access to these data sources and how their data is being used. The service providers and third parties who access the data should provide feasible tools to users to take advantage of these control options. We applied machine learning techniques to help predict privacy decisions a user/data subject is most likely to feel comfortable with. The prediction model can then assist users in making these decisions (e.g., reducing user burden and increasing the chance they end up with settings/choices that are aligned with their privacy preferences).

1.2 Tensions Between Privacy and Usability

Information privacy has become an increasingly challenging area over the past few years due to the variety of data being collected and the increasingly diverse ways in which it is used and shared across complex value chains. This explosion in the collection and use of users' data has created a fundamental tension between privacy and usability:

- From the perspective of the service providers, not all information platforms or entities offer sufficient or proper control to users. With less detailed "opt-in/opt-out" choices shown to them, users must adopt an "all-or-nothing" mechanism to be able to use the services.
- From the perspective of the users, even if the systems provide some level of notice and choice, users may lack knowledge of their viable options or find it challenging to manage the many privacy controls. Moreover, users most often do not read the privacy policies provided as the "notice" element of "notice and choice." Users also often lack time and motivation to look at the settings available to them, when such mechanisms exist. As Solove[86] pointed out, users may not fully take advantage of privacy control mechanisms if they are designed to be "self-managed" by the users. Indeed, a recent PEW survey[64] has found that 91% of people in the United States feel they have no control or not a lot of control over their data.

Developers and designers should take usability into consideration when designing systems that access private user data. Privacy by design, a concept popularized by Ann Cavoukian [24], contains seven principles for taking privacy into account when designing a system: (1) proactive, (2) privacy as the default setting, (3) privacy embedded into the design, (4) full functionality, (5) end-to-end security, (6) visibility and transparency, and (7) respect for user privacy. The concept advocates keeping privacy control user-centric. The privacy control should be proactive and protective by default. Privacy by design serves as a good principle for developers to provide sophisticated information services to users while simultaneously protecting users' information

privacy. Specifically, the developers or service providers should not use users' privacy-sensitive data in ways that the users are unaware of. Also, they should provide control for users to set their privacy preferences. In instances where users have diverse privacy preferences and difficulties in making decisions related to the settings, it is ethically important to assist users to configure the privacy settings to match their preferences, so that the users are less exposed to privacy risks.

There is tension between the level of control given to users over the collection of their information and the usability of the control made available to users. Nowadays primary web services, such as mobile phones, web browsers, and social networks, are providing an increasing number of privacy settings to users, which may give more control to users theoretically. While the number of privacy settings made available to users still falls short of capturing important considerations they care about, the number of settings has grown to become unmanageable (e.g.,[57]):

- **Privacy settings can be overwhelming.** Some privacy decisions might require users' intensive engagement. The service providers may provide only limited explanations of the outcome of each privacy decision. Therefore, users might not be well-equipped to make decisions when asked. Second, privacy is usually not a primary task of users when making privacy decisions. The presented privacy decisions might involve context-specific tradeoffs. Users might not have sufficient attention, time, or energy to spend on making each privacy decision.
- **Users can be affected by cognitive and behavioral biases when making privacy decisions.** Even if users are expected to have sufficient attention and patience to make privacy decisions, they may suffer from cognitive or behavioral biases, such as heuristic biases, anchoring effect, etc. These biases are discussed in more detail in Chapter 2. If users are not fully engaged and rational when making privacy decisions, they may commit errors or make regrettable decisions when configuring their privacy settings. Thus, it is reasonable to provide assistance to users and allow them to revisit their settings.
- **The number of privacy decisions is increasing. It is unrealistic for users to fully take advantage of all these privacy settings.** By enforcing awareness and control of privacy, the service providers also transfer some responsibilities of privacy protection to the end-users. As services become increasingly more sophisticated and users continue adopting these innovative services, the number of privacy settings users are exposed to also increases. It is important to note that users' privacy decisions are usually not black-or-white. Users are expected to make tradeoffs between the various benefits and risks presented within the specific context. Thus, it would be impractical to reduce the number of privacy decisions by enforcing universal default settings, especially since people have different privacy concerns and preferences in different contexts. As the default settings are usually a one-size-fits-all solution, users have to manage a massive amount of settings by themselves to align their privacy settings to their preferences.

1.3 Mobile App Privacy: A Typical Domain to Study

Mobile app privacy is a typical domain where we observe significant tensions between privacy and usability. Smartphones have been broadly adopted worldwide[27]. They have been evolved to have an increasing number of APIs to provide sophisticated functionality to users. Apps can access personal data to provide location services, social sharing, fitness tracking, and so on. Both Android and iOS use a permission mechanism to organize and control the access to various personal data collectible on the phones. Android introduced a permission control tool named App Ops in Android 4.3 (2012) and a refined version of permission manager in Android 6.0 (2015). For each app requesting a permission, users can choose "Allow" or "Deny" to manage the personal data access of that permission.

However, even though they have control over permissions, users are still facing an increasing number of permission settings to manage. On average, an Android app require five permissions [16]. And an average Android user installs more than 80 apps and open almost 40 apps every month[15]. So average users would need to manage over a hundred permission settings on their phones. The numerous Android permission settings users must manage are also a reflection of the diverse privacy preferences available to mobile app users when it comes to granting access to their data[61].

Our research focuses on reconciling privacy and usability on users' management of Android permission settings. The related work in this domain attempted to solve the following three aspects of this problem: improving interaction with users by increasing their awareness (such as [47]); alternative permission control mechanisms, such as context-based control[78]; and providing decision support to help users configure permission settings[12]. In the following, we demonstrate our design and study results regarding this issue:

- We conducted several studies to capture and analyze users' real app permission settings[60, 61]. LBE privacy guard[80] is a pioneering app permission management tool for early versions of Android phones. Based on an analysis of a collection of real Android users' permission settings from the LBE privacy guard app, we found that users' app privacy preferences are diverse. Also, we developed a permission manager app to collect users' actual app permission settings on their phones.
- We modeled users' diverse app privacy preferences and generated personalized recommendations for permission settings. Multiple previous approaches[12, 40, 43] proposed assisting users by displaying analytical results of privacy-sensitive behavior, and by providing decision support powered by expert labeling or crowd opinions. However, the recommended app privacy settings from previous methods were exceedingly generic. As we observed, users' mobile app privacy preferences are quite diverse. Therefore, we use a profile-based method to generate personalized recommendations for users in our permission manager app.

- We designed and implemented a personalized privacy assistant (PPA) for mobile app permissions[60]. The app serves as an easy-to-access portal for users to manage permission settings. The app offer assistance to users by first asking a few questions to users on their app privacy preferences, which is then used to estimate a privacy profile and display profile-based recommendation settings to users in the PPA app. The responses from the participants showed the potential for reducing user burden when managing app privacy settings.

1.4 Dissertation Overview

1.4.1 Research Questions

This dissertation primarily focuses on the burden of managing a large number of privacy settings made available to users. The core research question addressed in this dissertation is:

Can machine learning be used to build models of people's privacy preferences that can help them configure their privacy settings?

- Is it possible to build models that can help predict many of the app permission settings a given user would likely select?
- Are users open to accepting this type of technology? How should this technology be configured for the adoption of users?

1.4.2 Summary of Content

We combine machine learning techniques with human-subject experiments designed to evaluate the practical impact of our technologies in realistic contexts. In this regard, we collected and analyzed a large collection of privacy preferences, evaluated machine learning models and techniques, and examined the predictive power of different combinations of features. Moreover, we developed and piloted a personalized privacy assistant for Android users, which was tested during the course of 10 days with actual Android users using their phones during their regular everyday activities.

This dissertation consists of three main studies.

- The first study explores the use of machine learning in the context of a large corpus of permission settings obtained from actual LBE users. LBE is a rooted version of Android that had been allowing users to manually configure up to 12 different permission settings on an app-by-app basis for several years before the introduction of App Ops. The corpus to which we were given access included permission settings for 4.8 million users. We pre-

processed the data to only retain settings for users who had genuinely engaged with the permission settings. Thus, we obtained data for slightly more than 239,000 users. As part of this first study, we were able to analyze the predictive power of different combinations of features and evaluate both individualized learning models and cluster-based models. Our analysis reveals that even a small number of clusters can predict users' mobile app permission settings. As part of simulations run using our model, we were also able to evaluate semi-interactive scenarios in which machine learning is used to automatically configure settings based on our model's predictions. These simulations reveal, for instance, that if users are asked to specify only 10% of their settings manually, it is possible to achieve over 90% accuracy in predicting the other settings. These results suggest that such an interactive approach of combining machine learning with selective dialogues with users can help reduce user burden.

- In our second study, we focused on a smaller set of users to uncover how to elicit privacy preference data more efficiently and accurately. Indeed, the aim was to elicit preferences that also reflect the purpose for which different apps request different permissions – a dimension reported to be significant in prior research by Lin et al.[57, 59]. Here, we introduced the use of privacy nudges[10, 14] as a way of ensuring that users engaged with their privacy settings. Moreover, the privacy nudges were designed to include information about the purpose for which different permissions were used (e.g., an app's core functionality, sharing with advertising networks, sharing with analytics providers). The results further confirm that a small number of privacy profiles can often provide meaningful predictive power. We conducted a detailed analysis of these privacy profiles, which revealed people's mobile privacy preferences.
- In the third study, we proceeded to test our developed techniques on real users. We developed a profile-based personalized privacy assistant (PPA) app. The assistant asks users a small number of questions based on the actual apps they have on their smartphones to determine the closest privacy profile of the user. We studied a sample of 49 users who installed our personalized privacy assistant on their Android phones and used it as part of a 10-day trial period during their regular day-to-day activities. This pilot also included an additional 23 baseline users who were provided with the same functionality although without the benefit of the machine learning models. The results of this study suggest the success of our personalized privacy assistant at helping users configure their privacy settings: 78.7% of recommended settings were adopted by users and only a tiny fraction of these recommended settings were overturned at a later date in the study. After conducting the studies mentioned above, we modified the app design and deployed the app on the Google Play store for download by rooted Android users. Modifications made to the mobile app before its deployment in the Google Play Store are also discussed.

1.4.3 Contributions

In summary, this dissertation makes the following major contributions:

- We demonstrated that it is possible to build models of people's privacy preferences that can help predict many mobile app permission settings. We were the first to analyze a large-scale corpus of permission settings obtained from real Android users. We analyzed the predictive power of different combinations of features and machine learning techniques to predict users' permission settings. Our results show that with some of a user's settings and interactive responses from the user, we can build models to predict the user's permission settings with high accuracy.
- We found that while users' privacy preferences are diverse, a small number of clusters can go a long way in capturing and predicting people's app permission settings. We conducted a study to collect 84 Android users' app permission settings by using privacy nudges to motivate users to be engaged with the settings. We have shown that by using nudges, it is possible to obtain in-depth data on privacy preferences from a relatively small number of users and transform these preferences into actionable privacy profiles with strong predictive power.
- By conducting a field pilot study, we demonstrated that it is possible to design privacy assistants that exploit the power of machine learning to help people configure permission settings. We designed and developed an interactive privacy assistant app that interactively captures users' app privacy preferences and provides profile-based personalized recommendations for permission settings. Rather than automating privacy decisions, our privacy assistant recommends settings that users can review to decide whether or not to accept them. The results of our field study of this technology (N=72) suggest that this approach holds significant promise. People have reported finding this functionality to be helpful. They adopted many recommendations and appreciated the reduction in user burden. We also made several improvements to the app and made it available on the Google Play store.

The remainder of this dissertation is organized as follows. In Chapter 2 we discuss the key concepts and techniques, as well as the research conducted prior to or concurrently with our own work. Chapters 3, 4, and 5 each discuss a corresponding study. Chapter 6 discusses the Personalized Privacy Assistant app we designed and developed on Google Play store. Finally, we summarize the contributions, implications, and future work stemming from this dissertation in Chapter 7.

Chapter 2

Background and Related Work

Tensions between control and usability when it comes to adjusting privacy settings reflect a combination of complex challenges:

- Users' privacy preferences are diverse. People may have different levels of comfort or concerns over private data collection. End-users would still need to review and potentially adjust each individual privacy setting even though service providers might have applied one-size-fits-all default settings for all of them.
- Users' privacy preferences may depend on a number of contextual factors. As prior research has shown, people's privacy preferences are not just diverse but they are also complex and tend to depend on a number of "contextual attributes." For instance, a user's willingness to share his location with a particular app might depend on how the app will use the user's location, for how long it might retain this information, or where the user is at a particular point in time. This has been formalized by Nissenbaum and colleagues under a framework referred to as Contextual Integrity[69]. Contextual integrity introduces models of people's privacy expectations that are organized around five sets of parameters (e.g., data subject, sender of the data, data recipient, information type, and transmission principle). More recent research has produced finer taxonomies for some of these parameters in different domains (e.g., location sharing, mobile app privacy preferences, etc.[21, 32, 36, 47, 57, 59]).
- Users may have insufficient awareness of data collection, usage, and available options to control them. These complicating factors may make it more difficult for users to make good decisions. For example, if a user cannot know the exact third-party stakeholders who can access their data, it is difficult for this user to decide whether or not to allow the sharing of data to third parties.
- Users have bounded rationality when managing privacy settings. Users may be subject to disadvantages of cognitive and behavioral biases. For example, users may favor the short-

term reward of installing and playing with an app right away and discount long-term risks associated with granting sensitive permissions to this app. There are other biases such as heuristic biases, anchoring effects, and so on. It is important for researchers to understand these biases and design techniques to help users counter them.

In this chapter, we dig into these factors and review related work that directly or indirectly helps users in managing their privacy settings. Especially for the domain of mobile app privacy, we also discuss the work on improving the usability of mobile app privacy settings and helping users manage their app permission settings.

2.1 Segmenting and Modeling Users' Diverse Privacy Preferences

In this section, we discuss some representative work on capturing or modeling the diverse privacy preferences of users. There are generally two types of approaches: segmenting or grouping like-minded users, so that we can have nominal labels that describe users' preferences; or using a multidimensional feature space to describe a user's privacy preference.

2.1.1 Westin's Privacy Segmentation Index

Alan Westin is one of the pioneering researchers in the domain of consumer privacy, especially from the perspective of user perceptions and preferences. He conducted and advised on multiple privacy-related consumer surveys, which are collected and studied by Kumaraguru and Cranor[55].

Westin developed a segmentation index to categorize people according to their general privacy attitudes and discover the trends of the general public's privacy concerns. Segmentation was used by Westin in multiple different studies with different survey questionnaires and segment-assignment logic. The index typically consisted of three Likert survey questions and assigned participants based on their responses to the three questions in a rule-based manner. Each survey posed a different collection of questions, and Westin categorized the survey populations using different thresholds and methodologies. He showed that often when it comes to privacy, users can be organized into three broad categories[54, 55, 89]:

- Privacy Fundamentalists: The people in this group usually have high privacy concerns. They tend to feel that they have been invaded in privacy or have lost privacy control.
- Privacy Pragmatists: Those in this group usually have moderate privacy concerns. They make a tradeoff between benefits they can get from services and risks of giving their personal information to businesses.

- Privacy Unconcerned: People in this group have less anxiety about collection and use of their personal data.

This segmentation can be used as archetypes of user behavior when designing systems that access and use personal data. Similar segmentation can be developed for other domains that also require user awareness and control on managing settings or preferences. For example, Dupree et al.[31] categorized users according to their survey responses on their security-related attitude and behavior.

Westin's index has been widely compared with observed user attitude and behavior in other studies (such as [9, 25, 28, 29, 65, 97]). Here, we list some examples:

- Cranor and Reagle[29] (also[8]) conducted a survey in order to study internet users' attitudes about online privacy. The authors divided participants into three clusters, similar to Westin's segmentation, and then compared the clusters on several significant aspects in the survey responses. The authors also considered the differences and shared concerns of users across different clusters.
- Another survey conducted by Acquisti and Grossklags[9] asked participants about their privacy-related experience and attitudes. They categorized the users' diverse attitudes using clustering techniques to organize them into four groups: privacy fundamentalists with high concerns overall, two medium groups with focused concerns, and a group with low concerns in all fields.
- Chanchary and Chiasson[25] used Westin's index to divide participants in an online survey. They found that Westin's index is a significant factor affecting participants' willingness towards data sharing.

Westin provided qualitative explanations for each category of users' privacy attitudes. However, when it comes to understanding or predicting users' context-specific behaviors, researchers have raised concerns on the limitations of this segmentation[28, 48, 65]. The previous studies failed to observe significant correlations between the categories and the participants' intended or actual privacy-related behaviors. Woodruff et al.[97] argued that these simple segmentations have limitations in predicting heterogeneous context-specific privacy decisions in practice. Hoofnagle et al.[41] provided recommendations to improve the segmentation to adopt characteristics such as context and the differences of consumers as individuals.

In addition to segmentation, people have also applied linear scales on privacy attitude to describe their privacy preferences. The IUIPC Scale[66] and PCS Scale[23] are typical scales ranking users' information privacy attitudes in a specific internet domain context, using scenario-like questions instead of abstract ones[65]. However, similar to the Westin index, these scales provided linear or ordered representations of users' diverse privacy attitudes. While having an explanatory role in users' privacy mental model, these scales were not designed to have strong prediction power on users' intended or actual decisions on individual privacy settings or actions.

Knijnenburg et al.[52] found that the information disclosure behavior of users is rather multidimensional. Their findings indicate that, while good at explaining the phenomenon, these scales are not sufficient to model or predict users' tentative behavior or decision-making in a complicated context or environment.

2.1.2 Modeling Users' Privacy Preferences

In this dissertation, we are not directly trying to build models to predict users' measured privacy attitudes or concerns but to predict their tentative actions or decisions in specific scenarios or contexts.

- We are inspired by Westin's early exploration of segmentation and further adopted the clustering-based models to describe users' diverse privacy preferences using a relatively small number of privacy profiles.
- While the works cited above focused on capturing users' privacy attitudes or intentions, it is important to point out that these intentions are not directly mapped to users' actual behaviors such as their product choices or privacy settings[70, 87].

Beyond Westin's methodology, which constructs a segmentation in a heuristic way, we focus on further developing segmentation-based models that are optimized for predicting people's selection of actual privacy settings. These models can not only capture and explain users' diverse preferences but also assist users on tasks of managing their privacy settings.

Researchers have been working on modeling users' diverse privacy preferences from their self-reported or captured privacy decisions. One common focus would be looking for significant factors or characteristics affecting users' privacy decisions. Ackerman et al.[8] surveyed user concerns about online privacy. Using the method of factor analysis, they identified key factors in online information disclosure such as sharing data with third parties, identifiability, type of data, and purpose. Consolvo et al.[28] conducted a study on location disclosure preferences and found that the requester, granularity, and reasons for using location data are among the most important factors. The complexity of people's location sharing privacy preferences was also studied by Benisch et al.[21]. They quantified the benefits of exposing different types of privacy settings to users. Lin et al.[58] further studied the location sharing preferences across nations and key attributes of users' preferences. King et al.[50] studied users' privacy concerns using social networking sites. They found that instead of knowledge or expertise, the participants' experience in adverse privacy events is a better indicator of their tentative privacy concerns. Egelman et al.[32] found that the Big Five personality traits were a weak indicator of participants' privacy attitudes (measured by IUIPC[66] and PCS[23]). The solutions mentioned above provided useful insights into how users' decisions on privacy settings are formed. These methodologies are useful to identify key orthogonal factors of users' decision-making or to estimate correlations between each factors and users' choices of granting or denying access to personal data. However, these

solutions were not designed for the purpose of building models with high predictive power on users' individual privacy decisions.

In contrast to approaches that manually identify collections of factors to describe users' decision-making process, model-based methods can capture and predict users' decisions in a systematic manner. Benisch et al.[21] were the first to systematically quantify the different types of tradeoffs and compare the value of exposing privacy settings to users alongside additional contextual attributes (such as location, time of day, weekdays vs. weekend, etc.).

One direction of the model-based approach is to develop models with hidden factors. Factorized models may capture hidden or uninterpreted factors that affect users' decisions. Fang and LeFevre[35] applied this method and proposed a privacy wizard that tries to minimize user effort in configuring Facebook social sharing settings. The privacy wizard uses active labeling to trade the accuracy of captured privacy preferences against user burden. Wisniewski et al.[51, 96] conducted a survey on Facebook privacy behavior. They ran a mixture factor analysis and segmented users into six classes based on their responses in 11 dimensions using BIC criteria for estimating the number of classes.

Another approach is to develop privacy profiles (or personas) by clustering users with similar preferences together. This method makes it possible to develop privacy profiles that go beyond the scope of the three manually defined segments from Westin's Index.

- Locaccino[30, 95] is a location-sharing prototype to help users share their location with other people using fine-grained privacy control. It has a privacy wizard powered by a profile-based model[67, 79, 95] to assist users in configuring their settings for sharing their location data with different social groups.
- Peddinti et al.[75] analyzed users' perceptions of private data use in advertising. They grouped like-minded users into two clusters based on their survey responses and suggested providing a more detailed segmentation of users.
- Preibusch[76] conducted a similar study in which users were clustered based on their online privacy preferences over certain tradeoff items of privacy and functionality. Further, the author experimented with variance-ratio criterion to select the number of clusters and suggested favoring manageability of clusters over numerical variance measures.

Our work focuses on building clustering and classification models with predictive power to assist users in their privacy decision-making on mobile app permissions. Thus, instead of only grouping users to generate tight clusters or human-interpreted segmentation, we optimize our clustering model so that the learned profiles of like-minded users can better express their diverse preference space and provide higher prediction power for individual privacy decisions.

2.2 Privacy in Context

Users' privacy preferences for mobile apps is a complex space. General-purpose surveys such as the ones conducted by Westin focused on asking about users' privacy attitudes overall or on a particular domain. In contrast, when it comes to individual privacy-related decisions in their daily lives, such as granting data access to some party or enabling certain sensitive functionality on their devices, the context of the privacy decision plays an important role. For example, Schairer et al.[81] conducted a qualitative study on users' opinions on controlling their personal health information. They found that the participants' privacy-related behavior is contextual and habitual. In addition to the tradeoff of perceived risks and benefits, participants' behavior is also affected by their past experience, trust in the system, and so on.

Privacy preferences are closely tied to a number of different contextual factors. Instead of characterizing privacy violation as breaking the boundaries between privacy and public data or insufficient control, Helen Nissenbaum argued that the problematic privacy design is due to not following the users' expected norms in specific contexts. Moreover, Nissenbaum proposed a framework of "privacy as contextual integrity"[20, 69]. As a normative model, this framework is designed to evaluate if the information flows between entities are appropriate under certain circumstances. The contextual integrity is mainly determined by four aspects: "information norms," "appropriateness," "roles," and "principle of transmission."

The concept of privacy as contextual integrity can be applied to explain what end-users are facing in the management of their privacy settings. Nowadays, users are employing a vast amount of services in different contexts. In general, these different application scenarios have different norms. Also, users may have different levels of privacy concern, and thus could have quite different acceptance of flexibility in violating those norms. Further, results from Locaccino[21, 30, 95] have verified the theory of contextual integrity in the domain of social location sharing.

We find that mobile app privacy is a typical domain in which contextual integrity plays an essential part in privacy decision-making. Users install many apps on their phones for different purposes and uses. Also, even for similar purposes, apps behave differently and receive different levels of trust from users due to reasons such as popularity, usability, brand reputation, and so on[47]. Users' decisions regarding app permission settings could be affected by factors such as app behavior, the trust of users towards the app's developer or towards this category of app, the sensitivity of the data in the user's mindset, and norms about whether this kind of app needs the data and how it would be used.

Thus, theoretically, users would need access to fine-grained privacy settings on mobile apps. Also, the users would require diverse settings if they would like them to be consistent with their actual privacy preferences. Both Android and iOS now provide permission-based privacy control for each app on their devices. However, people's privacy preferences change in different contexts such as different apps, permissions, and purposes[57, 59, 60, 62]. King and Hall[49] also found

that mobile app users' comfort about private data access was directly related to the context of use.

Inspired by the idea of contextual integrity, in this dissertation we first analyzed users' mobile app permission settings from different contextual factors, such as the category of the app, the requested data type, the purpose of permission use, and so on. Then we used these contextual features to generate models of users' preferences that could be used to predict their app permission settings. As the mobile domain and the Internet of Things (IoT) expand, we could expect that more contextual signals will be available to end-users and to researchers. Thus, our framework could also be extended to adopt more contextual features in the prediction models for better accuracy and coverage.

2.3 Users' Privacy Decision-Making

2.3.1 Incomplete Information and Bounded Rationality

Configurations and management of privacy settings have become ubiquitously available to users in various services. As discussed above, these privacy settings are most often context-dependent. The process of privacy decision-making requires significant user involvement so that the actual settings are aligned with users' privacy preferences. However, users might not be capable of keeping them fully aligned. Users are shown to be affected or influenced by a number of cognitive or behavioral biases. Thus, even if the end-users are provided with sufficient awareness and control, it is still a non-trivial task for them to manage their privacy settings[9, 10].

End-users only receive incomplete and asymmetric information on privacy-related actions. For example, in the current iOS and Android ecosystem, apps may request access to privacy-sensitive permissions using simple explanations without formally defining why they need the access and how they will store and use the data. Users are expected to make privacy decisions, such as whether to install the app or whether to grant or revoke a permission, based on partial and sometimes unreliable information such as app descriptions, their trust in the developer, community reviews, their perception of app functionality, and so on.

Users have bounded rationality in privacy decision-making, even if they have sufficient information and knowledge to make privacy decisions. It could be difficult for them to effectively obtain and analyze information and react according to their preferences. The users who claim to be very privacy-concerned might not take actions to protect their private data when dealing with individual privacy settings. They may often be biased by heuristics, such as available examples they are able to obtain, familiarity with the brand, and availability of privacy policies. Users' privacy decisions are hampered by cognitive biases such as anchoring effect, loss aversion, framing, hyperbolic time discounting, overconfidence, post-completion error, and status quo bias[9]. Also, the users' perception of their privacy protection may not be objective. Both objective or

framing changes may impact their decision-making[11].

Users may be affected by hyperbolic discounting when making privacy decisions. For example, users may be excited to try a new app or get distracted from work or entertainment when they receive pop-ups to configure permission settings. They may discount the value of the later reward, which is privacy protection that aligns with their preferences. Later, if users are surprised at how their data has been accessed or if they are informed by privacy-related notifications, they have to invest more effort to revisit and adjust their settings to stop unintended data access or leakage.

As analyzed, users have various disadvantages in the task of managing their privacy settings. The phenomenon eventually leads to regrettable privacy settings for users and their dissatisfaction with the services due to privacy concerns.

Privacy is typically a secondary task for users, particularly mobile app users. In older versions of Android, users are shown privacy information of apps only during installation, when they favor short-term benefits over long-term risks. Felt et al.[37] found that only a small fraction of Android users paid sufficient attention to managing Android permission settings in an early version of Android. The users suffered from insufficient knowledge or information on mobile app permissions when making decisions. Also, cognitive overload affects users' decision-making with respect to mobile app privacy. The authors suggested reducing the number of permission settings or warnings by only showing the settings with high risks. Considering that users' perceived risks are different even on the same app accessing the same permission, our Personalized Privacy Assistant (PPA) in this dissertation did not provide universal default settings and hide the low-risk ones but offered personalized recommendations in order to help reduce user burden. Another adverse effect if we assign full burden to users to manage their app permission settings is that they could become fatigued. Korff and Bhme[53] found that participants had more negative feelings and regrettable decisions and less satisfaction when exposed to a significant number of privacy options.

It is encouraging to see that Android has reorganized the permission mechanism since version 6.0 to group similar permissions and simplify the user interface with the app permission manager. However, this change did not solve the fundamental problem. Users will expect to have more powerful smartphones with enriched functionality that collect more types of personal data and an increasing number of apps accessing these APIs. Thus, they are still facing an increasing huge burden of managing their privacy settings on their phones.

2.3.2 Privacy Nudging

As discussed above, prior research has in part focused on organizing contextual dimensions that influence people's privacy preferences and has also looked at segmenting populations of data subjects based on how they feel or act under different conditions. Yet another line of research

has focused on motivating users to pay more attention to privacy decisions they have to make, and help them make better-informed decisions. A lot of this research has revolved around the development and evaluation of so-called "privacy nudges."

Privacy nudging[10] is a soft-paternalistic mechanism to help users by informing them and guiding them towards more protective privacy choices or decisions without imposing a specific course of action. Soft-paternalistic interventions attempt to influence users' decisions while at the same time provide freedom of choice, advocated by libertarian solutions.

To design a specific privacy nudge notice, one needs to consider the following contexts[10]:

- Information: reducing hurdles such as information asymmetry.
- Presentation (framing, structure, and so on.): reducing psychological or behavioral biases such as anchoring and bounded rationality.
- Defaults: reducing status quo bias and diversification heuristic.
- Incentives: increasing users' motivation to behave according to their actual preferences by means of, for example, rewards and punishment.
- Reversibility: reducing the impact of mistakes or errors.
- Timing: defining or optimizing the most appropriate time to display the nudges.

Schaub et al.[82] provided systematic guidance for researchers and developers to design effective privacy notices, which are mostly applicable to the design of privacy nudge notices. They proposed four dimensions to design a privacy notice: Timing, Channel, Modality, and Control. In addition to the nudge notice dimensions mentioned in the list above, the notices should also come with easy-to-access control so that users may perform immediate reactions to these privacy notices.

Specifically, for the domain of mobile app privacy, considering the system resources available in the current Android and iOS ecosystem, nudges could be designed to include more detailed and organized information about permission requests and alert the users for necessary actions on configuring their permission settings. Multiple studies have provided evidence suggesting that the success of such nudge notices can be sensitive to the design, such as content and timing. Felt et al.[36] conducted an online survey and found that participants differentiated between potential risks of mobile app data access and felt annoyed with low-ranked risks. Their work indicates that the nudges are more effective when they are minimally interruptive and personalized for individual concerns. Balebako et al.[19] conducted a study on the timing of mobile app privacy notices. The results indicated that showing notices while using the app can lead to significantly higher recall from users, compared with showing notices right before app installation. Patil et al.[73] suggested that for location sharing, it is better to provide actionable notices as runtime feedback or to delay the notice to avoid unconscious reactions from the users.

Hazim Almuhimedi studied the design of privacy nudges concerning app privacy for Android

users. In his paper[14], he conducted a field study on showing privacy nudge notices to Android users. The results showed that the daily nudges motivated participants more effectively in terms of revisiting the permission settings and restricting the permissions that they otherwise ignored. In his later dissertation[13], Almuhimedi also evaluated a number of different ways of framing nudges such as the context (location, back-end or front-end, etc.), usage/purpose, and potential outcome (inference of places of interest, etc.). He also conducted a first study on habituation issues with privacy nudges. He found that users' engagement decreases when they receive repeated nudges. In our work, beyond using nudges as "soft-paternalistic" interventions, we also used them to motivate users to engage, review, and adjust the app permission settings in the lab study environment, so that the settings collected from participants might better align with their actual preferences for apps on their phones.

In this dissertation, inspired by the idea of privacy nudging, we designed a privacy nudge message to help us better capture and model users' permission decisions in the field study. Different from the use of nudges in the related work above, we have two purposes in applying privacy nudges:

- Considering the situation that the field study we conducted has a limited length of time available for each participant, it is necessary to motivate participants to interact with our app and review their settings so that the settings on their Android phone align well with their privacy preferences. The nudge notifications helped the participants to contribute their settings in a relatively short period. We treat the settings users committed to during the study after daily privacy nudges as the ground truth of users' app privacy preferences.
- When it comes to receiving personalized recommendations for permission settings, it is possible users could agree to or accept the recommendations with incomplete involvement or consideration. To further confirm that they are comfortable with keeping these permissions on the phone, we show daily nudge notifications about the facts of permission use after users adopt the recommended setting changes. These nudge notifications could help motivate users to review their settings and adjust them if they find some regrettable acceptances.

2.4 Mobile App Privacy

Mobile app privacy is one common domain where the tradeoff between privacy and usability affects the effectiveness of privacy control. Android and iOS, the two major ecosystems, are increasingly providing more system APIs for app developers to build more sophisticated app functionality. However, the APIs allow the apps to collect many types of privacy-sensitive data that may raise user concerns[84]. Both platforms use permission-like systems to organize these API accesses. Both Android and iOS have adopted this mechanism to provide a privacy control tool for users to grant or deny permissions requested by apps.

- Starting from iOS 6 in 2012, Apple has introduced privacy settings for apps on Apple devices. For each permission requested by an app, a user can toggle the switch to turn the permission on or off. In 2014, Apple introduced finer granularity on location control so that users could choose to allow location to be always accessed or to be accessed only while using the specific app in the foreground.
- Android has adopted a permission mechanism since the earliest versions of the mobile operating system. However, prior to Android 4.3 Jelly Bean (2013), users were only shown a list of accessing permissions for a specific app before installing the app in the app store. The user could proceed by clicking the "install" button, thereby implicating that they allowed all permissions requested by the app. Starting with Android 4.3, Android introduced a hidden app permission manager named App Ops. By default, users were not able to access the settings. This app was only accessible if users installed a third-party tool to trigger a specific Android Activity to be launched. In Android 4.4.3 Kitkat (2014), Android kept expanding the permission model to allow more resources to be controllable by App Ops but disabled the app from being accessed. In Android 6.0 Marshmallow (2015), Android introduced a brand-new permission manager for Android apps. The permission manager simplified the permission control and display by grouping permissions into groups. Also, the permissions by apps could be requested in a runtime manner. When an app requested permission for the first time, Android would pop up a runtime dialog for users to configure whether to allow or deny the specific permission request.

2.4.1 Better Awareness and Control of Mobile App Privacy

To demonstrate sufficient and easy-to-comprehend information about privacy actions, risks, and potential consequences, it is crucial to design a usable user interface with mobile app permission settings. Kelley et al. found that users are not well prepared to handle these mobile app permission settings[45, 46, 47]. They[45] proposed a "nutrition label"-style approach to visualize information about permission requests for mobile apps during installation, to better help users locate the information they need when making app installation decisions. Balebako et al.[18] conducted a field study on visualizing private data leaks on real Android phones and found that participants did not expect certain private data sharing by apps. Fu et al. made a similar proposal to include disclosure on access frequencies[38]. Choe et al.[26] suggested using positive visual framing to guide users towards more protective decisions with respect to privacy when choosing apps to install in the app store. Harbach et al.[39] experimented with a design for an Android app installation screen where personal examples were added to better communicate potential risks of apps accessing specific permissions. They found that participants were guided toward protective decisions and paid more attention to the management task. These solutions helped provide guidance and future potential in terms of enhancing awareness and attention for users to make better privacy decisions.

System-side innovations have been proposed to automatically analyze mobile apps for potential privacy-related behavior such as data access and flow, risks, and consequences. TaintDroid[33], developed by Enck et al., enabled taint tracking of privacy-sensitive data by modifying the Android system into a custom ROM. Researchers could further detect privacy leaks and track real-time privacy-sensitive actions on Android phones. PrivacyGrade[59] uses static analysis on the Android apk files and crowd-sourced data labeling to provide privacy ratings and detailed explanations of the potential purposes of use, for permission accesses on Android apps. Multiple papers (such as[63, 99, 100]) have proposed to use both static code analysis and text analysis techniques to identify potential violations of privacy policies and regulations by apps. Styx[17], proposed by Bal et al., is a framework to express and demonstrate the potential privacy risks of Android apps. The system analyzed information flow of sensitive data and potential consequences so that it could help users better comprehend and compare different apps regarding their privacy risks.

Researchers have proposed several approaches to effectively redesign the protocol of controlling mobile app privacy settings. Evolved from AdDroid[74], PEDAL[62] implements bytecode instrumentations to allow Android users to separately control permission access used by app functionality and by advertising libraries. Apex[68] enables finer granularity of controlling location data access by mobile apps. Rahmati et al.[78] suggested adding more context-specific factors in the control. However, these advanced solutions did not consider reducing users' cognitive or physical burden in handling the settings. Restricting permission access by Android apps could result in loss of app functionality or comfort using the app. One workaround would be to mock the API data instead of denying the API access. MockDroid[22] implemented this functionality by modifying the Android system code of PackageManager. This approach may resolve app crashing or service denial by apps on a temporary basis. The app crash problem could be gradually solved in the Android and iOS ecosystems by enforcing app developers to correctly and smoothly handle user denials of permission requests.

Our work in this dissertation is also designed to improve users' awareness and control of mobile app privacy. We developed a modified version of the Android permission manager app in a field study described in Chapter 4. The app showed additional statistics regarding permission use and allowed easy-to-access switches for permission settings. We also developed a PPA app that additionally provided interactive dialogs and personalized recommendations on app permission settings to reduce user burden. Our methodologies could be integrated and extended with future scenarios of mobile privacy management, such as the proposed mechanisms cited above in this section.

2.4.2 Helping Users Configure Mobile App Permission Settings

Beyond improving the interactive scenario and protocol of manually configuring app permission settings, researchers have proposed multiple strategies to provide automatic decision support for

users on this task.

- Henne et al.[40] conducted a focus-group study on configuring location access of Android apps in finer granularity. They provided crowd-powered recommendations on the settings but suggested that the recommendations could be improved if differences in user preferences were taken into consideration.
- Ismail et al.[43] proposed that one could derive an appropriate tradeoff of privacy and usability by crowd-powered ratings of an app under different conditions of acceptance or denial of permissions. However, the authors did not diversify the recommendations to different users. Also, they did not check whether users would agree on the ratings of the same configurations, in their small-scale study (N=26).
- ProtectMyPrivacy[12] is an app privacy manager for early versions of jail-broken iPhones. It provides recommendations for privacy settings by aggregating the popular majority decision if a certain threshold is reached. The recommendations have been mostly accepted by users (67.1%) and could be generated from a small fraction of expert users. XPrivacy [1] implemented similar recommendations for Android apps using popular choices.

The solutions proposed above focused on providing universal recommendations based on popular choices overall. However, considering that users' app privacy preferences are diverse, personalized recommendations may help converge the app permission settings more quickly to match the actual app privacy preferences of users.

We are the first to design and implement a personalized assistant for app permission settings[60]. We have published the tool on the Google Play Store since 2017.[2]

- Inspired by various segmentation efforts on users' privacy attitudes and behavior, we developed a personalized privacy assistant that provides recommendations based on a crowd-powered engine that treats users differently according to their perceived preferences.
- One advantage to the non-personalized solution is that users can immediately receive recommendations without many detailed configurations or a "cold-start" period for the machine learning models to learn users' preferences. To reconcile usability and privacy, we developed a profile-based method so that users only need to answer a small number of questions prior to receiving personalized recommendations for app privacy settings.
- Due to the resource limitations of research studies, we applied privacy nudging to encourage more engagement of users on the management of permission settings. We implemented and experimented with the tool with real Android users in a field study, with recommendations and other settings actually being applied to participants' devices during the study.

[1] https://www.xprivacy.eu/

[2] https://play.google.com/store/apps/details?id=edu.cmu.mcom.ppa

Chapter 3

Can We Predict People's App Permission Settings? A Large Corpus Study

3.1 Introduction

As we discussed in Chapters 1 and 2, mobile app privacy is a good example of an application domain where users are presented with an overwhelming number of privacy settings. Smartphone users face a burden of overwhelming app permission settings:

- Mobile app permissions cover diverse sets of situations, some perceived as potentially being riskier than others by some users. For example, Google Map may access users' location data to provide navigating directions; however, a social app may use users' location to share their position with their friends and find nearby users. Thus, under the current Android and iOS mechanism, users would need to make decisions case by case.
- In addition, not all users have the same concerns even when considering identical situations. It is necessary for us to understand the diversity of users' app privacy preferences and seek potential ways to better set up the default privacy settings tailored to users' needs.

In this chapter,[1] we report our analysis result from a large corpus of permission settings from real Android phone users using LBE privacy guard. Powered by this dataset, we are the first to study users' app permission settings on a large scale. We analyzed the diversity of users' preferences for mobile app permissions and studied the potential adoption of personalized default settings to simplify the privacy decisions users are expected to make. We obtained a dataset of permission settings from 239 thousand users who actively managed their permission settings, out of the 4.8 million user base.

Our research goal is to study whether we can use machine learning to help users configure

[1]The work in this chapter was also published in[61].

their mobile app permission settings. In order to better understand how users manage these app privacy settings, it is crucial for us to get a closer look at smartphone users' app privacy behavior and preferences on a large scale. By analyzing the real users' app permission decisions, we can understand the diversity of users' privacy preferences for mobile app permissions and explore solutions that could potentially help users in this task of app privacy settings management.

In this chapter, we show that even though users' app permission settings are diverse, it is possible to build models to predict users' permission settings with a relatively small amount of input from users. By analyzing the diversity of the settings and trying different feature combinations of classifiers, we found that by selectively prompting users to provide input on items with lower confidence of prediction, we were able to boost the prediction accuracy. We also found that it is possible to define a relatively small number of profiles that would enable us to capture and simplify many of the permission decisions users are expected to make.

3.2 Permission Settings Dataset

3.2.1 LBE Privacy Guard

We obtained the dataset from LBE Privacy Guard,[2] a pioneering Android permission management tool. This is a privacy and security app that requires a rooted Android phone and allows users to selectively control the permissions they are willing to grant to apps on their phones[80].

Running on earlier versions of Android in the year 2013, LBE Privacy Guard relies on API interception to give users the ability to review up to 12 permissions that can possibly be requested by an app: "Send SMS," "Phone Call," "Phone State," "Call Monitoring," "SMS DB," "Contact," "Call Logs," "Positioning," "Phone ID," "3G Network," "Wi-Fi Network," and "ROOT." The nature of these permissions is very similar to that found in canonical versions of Android.

For each permission request from an app, the LBE app has four different possible settings for users to choose (also see details in Figure 3.1):

- "Allow": The user grants the app access to the permission.
- "Deny": The user denies the app access to the permission.
- "Ask": Each time the app calls the corresponding APIs, the system pops up a window prompting the user for a one-off decision. The window follows a 20-second countdown. In the absence of a decision within 20 seconds, the system assumes a "Deny" response. Users can also check a "Remember my choice" box to indicate that they would like their decision to become permanent (until they possibly change their mind). In this case, the settings remembered by the system change from "Ask" to either "Allow" or "Deny" (see

[2]LBE Privacy Guard https://forum.xda-developers.com/showthread.php?t=2320843

Figure 3.1(d)).

- "Default": This indicates that the user has never manually modified the settings. Default settings are interpreted according to the following logic: "Allow" if the permission requested is among "3G Network," "Wi-Fi Network," or "Phone ID"; "Allow" if the corresponding app is in a list of "trusted" apps; and "Ask" in all other cases.

After every new installation of an app on the phone, LBE will notify users to specify their permission settings for this app. Users can at any time revisit these permissions and elect to modify their settings for a given app.

LBE Privacy Guard is distributed on Google Play Store as well as several third-party app markets for rooted Android devices. It is also pre-shipped with a customized Android ROM called MIUI,[3] which is fairly popular in China. We obtained and analyzed a dataset that captures the permission settings of a total of 4.8 million LBE Privacy Guard users.

3.2.2 Permission Log Data of LBE Privacy Guard Users

Our dataset comprises the permission settings of 4.8 million LBE users in the form of permission logs collected over 10 days – from May 1, 2013 to May 10, 2013. Each log record contains permission settings for all the apps (identified by package name) installed on a given device. For each app, the log records the list of permissions the app requests and the most recent settings for these permissions (namely "Allow," "Deny," "Ask," or "Default"). Each device has a unique hashed device ID. The term "user" here refers to a unique Android device running the LBE app. For our analysis, we assume that each Android device corresponds to a distinct user. Apps are packages and are also represented by unique IDs. Our dataset does not include app information such as installation files, versions, or app store from which an app was downloaded. The LBE app is always running on the phone either in the front view or in the background. It periodically detects if a Wi-Fi network is available. If so, the app tries to upload its log. At most one log is uploaded each day. The logs are sent regardless of the operational status of the app. If the app is not running in ROOT mode or not functioning properly, the log will include "Default" for all the app permissions. Below we discuss how we sanitized our dataset to deal with these types of issues.

Over the 10 days, the dataset collected information about 4,807,884 unique users and 501,387 unique apps. The dataset comprises a total of 159,726,054 records, with a total of 118,321,621 unique triples of the form [user, app, permission]. It is worthwhile noting that, among the 4.8 million users in the dataset, 159,011 (or 3.4%) modified their settings for at least one app-permission pair over the 10-day interval. Among them, 2,978 (0.06% of the users) went back and forth for at least one permission setting. In our analysis, we focus on the final settings collected for each user over the 10-day interval. In other words, we do not limit ourselves to those users who modified

[3]MIUI: http://en.miui.com/

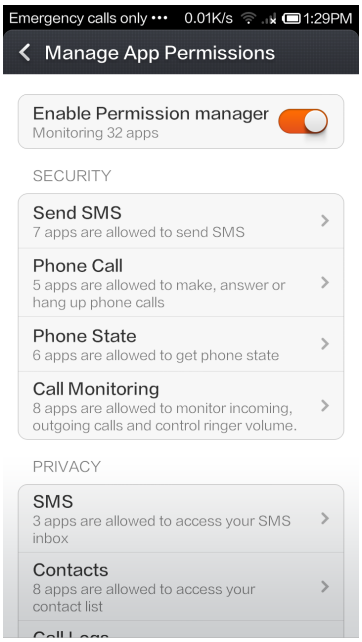

(a) The home screen where users can start managing app permission settings.

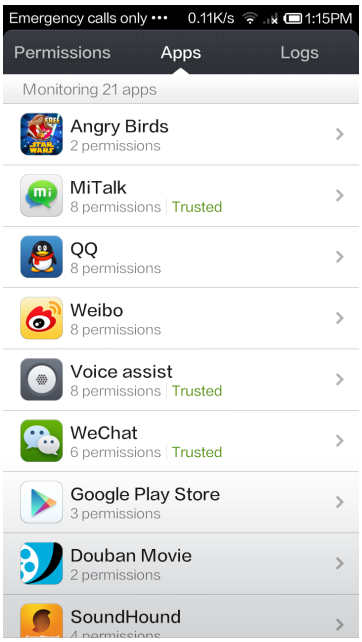

(b) The screen that shows the list of apps installed on the device monitored by LBE. For each app, the screen also shows the number of permissions it requested.

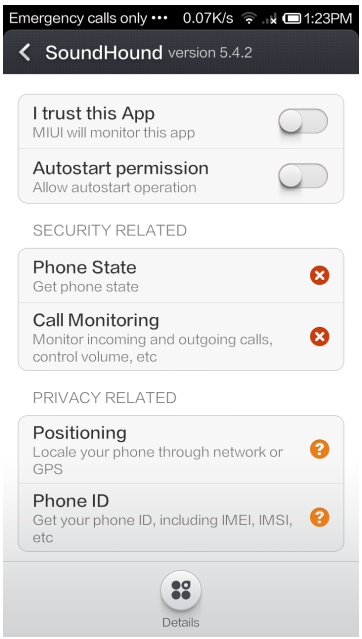

(c) The screen to manage the permission settings of a specific app. The example shown here is an app named SoundHound. Users can toggle the colored icons to change the permission settings.

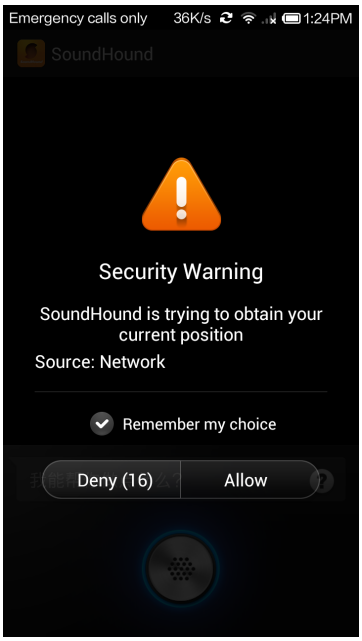

(d) If a permission is set to "Ask," this is a sample screen where the corresponding app is requesting the permission. The system pops up a dialog and asks for users' decision.

Figure 3.1: The App Permission Control of LBE Privacy Guard App on a MIUI 2S Phone

their settings during the course of the 10 days. This is further discussed below.

3.2.3 Data Preprocessing

Because our objective is to study people's privacy preferences as they pertain to the 12 permissions captured in the dataset, we proceeded to remove entries that might bias our analysis. In particular, we decided to focus on users who had actively engaged with the permission settings. This is in contrast to users who passively accepted them, or downloaded the app on a phone that was not rooted (in which case the user cannot control the settings), or perhaps did not even realize they had the ability to manipulate the settings. In addition, we also decided to focus on mainstream apps and removed entries that may correspond to more esoteric ones such as apps found only on secondary app markets. This is further detailed below.

- We limit our analysis to users who (1) have installed at least 20 apps requesting at least one permission, and (2) have manually selected at least one "Deny" or "Ask" setting for a permission request. These restrictions are intended to eliminate users who have an unusually low number of apps on their phones, and users who for one reason or another did not engage with the permission settings.
- We limit our analysis to mobile apps that (1) have at least one permission request, (2) have at least 10 users in our dataset, and (3) were available on the Google Play store over the 10-day interval of this study. This last requirement is intended to limit our analysis to mainstream apps, in contrast to apps from less reputable stores, which might prompt users to adopt more cautious settings and possibly distort our analysis.
- Finally, we also removed app permissions that were only recorded for five or fewer users. These app permissions are assumed to correspond to exotic versions or modified versions of some apps.

After this preprocessing, our resulting dataset still had a total of 239,402 users (5.0% of the initial population) and 12,119 apps (2.4% of the initial count). The number of decision records for these users and apps totaled 28,630,179 (or 24.2% of all records we started with). On average each user had 22.66 apps on his or her smartphone. This preprocessed dataset was deemed sufficiently large and diverse to warrant meaningful analysis, without being subject to the possible biases discussed above.

3.3 Data Analysis

3.3.1 Diversity of Users' Preferences

As already indicated, each representative user in our dataset had an average of 22.66 representative apps. On average a random pair of users had 3.19 apps in common, and each app requests an average of 3.03 permissions. A high-level analysis of user settings for different app-permission pairs shows that while there are some app-permission pairs on which the majority of users agree, there are also many such pairs for which users have diverging preferences. For instance, if one considers permissions for the top 100 apps, users agree on settings for only 63.9% of the app-permission pairs associated with these apps, if agreement is defined as 80% or more of the users selecting the same settings for a given app-permission pair (e.g., granting Angry Bird access to one's location). If one considers all the app-permission pairs for which we have at least five users, 80% agreement drops to 51.4%.

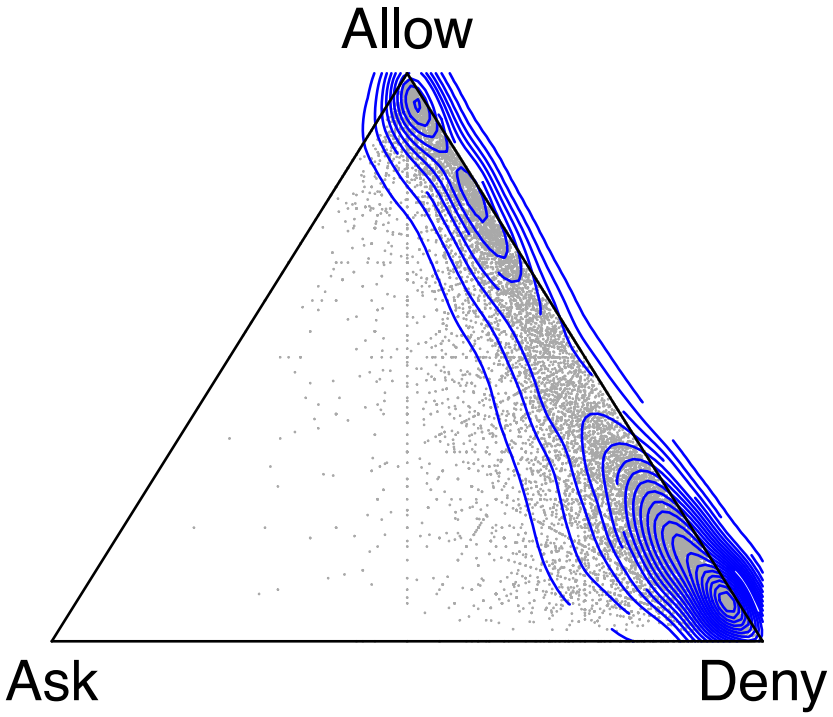

Figure 3.2: Distribution of Users' Permission Decisions for Each App-Permission Pair

Figure 3.2 plots the density and contour curves of app-permission pairs with at least 10 user decisions based on the mix of decisions recorded for each of these pairs. Specifically, the top corner corresponds to a mix where 100% of users "allow" an app permission, the bottom left corresponds to the case where 100% of users "ask" to be prompted for an app permission, and

the bottom right a mix where 100% of users select "deny." While many dots, each representing an app-permission pair, are concentrated around the top and bottom right corners, many are not (e.g., dots concentrated along the right side of the triangle). The plot also shows an overall bias towards either granting permissions or denying, with few users requesting to be prompted.

From the results above, we observed that most of the users did not settle into the "always ask" option, which could be even more burdensome for users to manage app permission settings. A large considerable amount of permission requests from apps did not have agreement on user decisions. This indicated that users had different levels of privacy concerns and context-specific preferences. With personalized learning, we could potentially provide personalized default settings that better fit individual users' preferences.

3.3.2 Modeling and Predicting Users' Decisions

As discussed above, we observed that in the dataset, users had diverse preferences on the vast amount of app permission settings they need to manage. With users having an average of 22.66 apps each and each app requesting 3.03 permissions, users are theoretically responsible for manually making more than 60 privacy decisions. An obvious question is whether this number of privacy decisions could be reduced by automatically predicting the settings a user would want to select – recognizing that not all users feel the same way and that therefore a one-size-fits-all model is unlikely to work. We argued that instead of providing a universal default permission settings for all users, it is crucial to provide users with more personalized decision support for permission settings, such as personalized default settings, or an interactive wizard or assistant tool for users to manage permission settings in a personalized way.

One basis for providing personalized default settings or assistance would be capturing users' preferences and predicting their tentative preferred decisions when they need to configure settings for new permission requests from the apps. Thus, we first experimented with whether we could predict users' app permission settings. Given that our main motivation is to alleviate user burden, we limit ourselves to a model where the set of decisions is restricted to "Allow" or "Deny," i.e., we exclude the "ask" option. Specifically, we look at whether it might be possible to build a classifier that could be used to predict a user's app permission setting in the form of a function:

$$f : (user, app, permission) \rightarrow decision$$

The prediction model is trained using a collection of decision records in the form of user, app, permission, decision quadruples. As we further trim our dataset to limit ourselves to decisions that are either "Allow" or "Deny," we are left with a corpus of 14.5 million records corresponding to a total of about 239,000 users and 12,000 apps. Through experimentation, we have found that good results can be obtained by using a linear kernel SVM as our model. This model also has the advantage of being quite efficient computationally[98]. The results reported below were obtained

using a toolbox called LibLinear[34] with both L2-loss dual support vector classification with linear kernel and L2-loss dual logistic regression to train the classifier with highest prediction power under linear kernel complexity.

Below, we report results obtained using 10-fold cross-validation. It is important to note that we did not treat all app permission setting data points uniformly and shuffle them for predictions. Instead, we organized them by users to simulate realistic cases in which the system obtained a subset of a user's settings and sought to predict the user's settings that were not yet observed.

- We randomly split all users into 10 groups of equal size.
- For each fold, one of the 10 groups is used for testing and the other nine groups for training. For each user in the training set, all the decision records (Allow and Deny) of this user are used to train the classifier.
- For each user in the test group, we randomly choose 20% of the apps installed by the user and the corresponding permission decisions made by the user (Allow or Deny) for training as well. This data could be obtained by looking at apps already installed by the user or by just asking the user to make some decisions for a small group of randomly selected apps – equivalent to asking the user a few questions.
- The remaining 80% of the apps downloaded by users in the test group are used to evaluate the accuracy of the classifier.

One challenge with using our dataset has to do with its high dimensionality coupled with the sparsity of data: a typical user has a little over 20 apps, but the dataset contains over 12,000 apps. We experimented with a typical technique for overcoming this challenge by using Singular Value Decomposition (SVD) to impute the unobserved values. The technique can produce a more compact, yet essentially similar dataset by effectively projecting the data along with a limited number of eigenvectors that collectively capture most of the information contained in the original dataset. To this end, we define a preference matrix of preferences P, where each entry in the matrix corresponds to the decision of a user (u) for a given app-permission pair (m). Specifically:

$$P[u][m] = \begin{cases} 1, \text{if the user chose "Allow" for app-permission pair m} \\ -1, \text{if the user chose "Deny" for app-permission pair m} \\ 0, \text{if no selection has been recorded} \end{cases}$$

To the extent that many users share similar preferences, one can expect the rank of this matrix P to be much smaller than either the number of users or the number of app permissions. In our analysis, we used the "irlba" toolbox[56] in R and its implementation of the SVD algorithm to produce a more compact dataset. The SVD method transforms the matrix P as:

$$P = U \cdot \Sigma \cdot t(V),$$

where $U \cdot t(U) = V \cdot t(V) = I$. Σ is a $u \times m$ diagonal matrix of eigenvalues, which are sorted in descending order. The SVD algorithm directly calculates an N-rank approximation of matrix P as the N-rank truncation of the initial matrix form:

$$P \rightarrow U' \cdot \Sigma' \cdot t(V')$$

We then generate the feature vectors of users and app permission pairs as follows:

$$F_U = U' \cdot sqrt(\Sigma') F_M = V' \cdot sqrt(\Sigma')$$

where $sqrt(\sigma')$ is a diagonal matrix whose diagonal values are the square roots of the corresponding diagonal values in Σ'. For each user u and app permission pair m, we then have:

$$P[u][m] \rightarrow P'[u][m] = F_U[u] \cdot t(F_M[m])$$

In the analysis, we limited the dimensionality to the 100 most significant eigenvectors (N=100).

An alternative to using SVD involves simply aggregating each user's settings by permission. This can be done using the matrix of preferences P, where each entry in the matrix aggregates decisions made by a given user u for the corresponding permission p, as:

$$P[u][p] = \begin{cases} \frac{a-d}{a+d}, \text{if user u chose "Allow" a times and "Deny" d times for permission p} \\ 0, \text{if no selection has been recorded for user u on permission p} \end{cases}$$

Thus the matrix cells have a value range of $[-1, 1]$, indicating the different levels of tentative decision of users on the specific permission, from mostly deny to mostly allow.

3.3.3 Performance of the Default Settings Prediction

Preliminary analysis suggests that people's privacy preferences when it comes to granting permissions are diverse. In this subsection, we take a closer look at the importance of different features in building classifiers that can be used to predict a user's permission decisions. We use 10-fold cross-validation. Also, we include in the training set permission decisions for a random 20% of the apps installed by users in the testing group. This is intended to capture scenarios where we use privacy preferences for apps a user has already installed, to predict permission decisions for new apps the user downloads on his/her phone.

Table 3.1 summarizes the 10 feature sets considered in this particular part of our study. They include a feature set where we aggregate decisions across all users and all apps (FS-1); a feature set where we aggregate decisions across all users and all permissions (FS-2); one where we aggregate decisions across all users for each app permission (FS-3); one where we aggregate decisions for each user across all apps and all permissions (FS-4); one where data is organized by user ID and permission ID (i.e., aggregated across all apps for each user-permission pair)

Table 3.1: Cross-validated Accuracies of Different Feature Compositions

	Features	Accuracy (%)	Note
FS-1	Permissions	67.13	Modeling from users' preferences overall
FS-2	App IDs	64.28	
FS-3	Permissions, App IDs	77.67	
FS-4	User IDs	71.48	Modeling using users' individual preferences
FS-5	User IDs, Permissions	80.72	
FS-6	User IDs, App IDs	76.13	
FS-7	User IDs, Permissions, App IDs	85.03	
FS-8	FS-7 + Aggregated user preferences on each permission: P[u][p]	**87.80**	Estimation from users' aggregated preferences
FS-9	FS-7 + P'[u][m] from SVD of users' decisions on top-200 apps	80.95	SVD estimation from the matrix of users' decisions on popular app-permission pairs
FS-10	FS-7 + P'[u][m] from SVD of users' decisions on top-1000 apps	80.66	

(FS-5); one where data is aggregated across all permissions for each app-user pair (FS-6); and one where data is broken down for each user by app-permission pair (i.e., user-app-permission triples) (FS-7). We also consider three feature sets where FS-7 is enriched with:

- The 12-permission user profiles, referred to as Feature Set 8 (or FS-8) in Table 3.1.
- An SVD model of user-permissions obtained by focusing on the 200 most popular apps in the dataset.
- An SVD model of user-permissions obtained by focusing on the 1,000 most popular apps in the dataset.

As can be seen in Table 3.1, looking at the prediction accuracy obtained with each of these feature sets, users, apps, and permissions all contribute to enriching the model and increasing its predictive power, with FS-7 (accuracy of 85.03% and StdErr = 0.08%) outperforming the other six feature sets FS-1 through FS-6. Supplementing these features with SVD models based on the top 200 or 1,000 most popular apps does not help and in fact results in lower predictive accuracy. On the hand adding user profiles based on the 12 permissions (FS-8) does enhance accuracy, bringing it from 85.03% to 87.8% (StdErr = 0.06%). The lack of improvement with the SVD model could be due to the fact that we took too many apps into account (200 and 1,000 most popular apps). A model based on a smaller number of apps (which would increase the likelihood that many users share a more significant fraction of the apps) could yield better results. The improvement based on the 12-permission model suggests that simple profiles based on aggregating user decisions along each of the 12 permissions provide additional discriminative power. Intuitively, this amounts to differentiating between different groups of users who may

be more or less comfortable granting different combinations of permissions across many apps. (e.g., people who have a problem disclosing their location versus people who do not mind).

3.3.4 Evaluating Simulated Interactive Scenarios

While 87.8% accuracy is promising, it is easy to imagine that even higher accuracy could be achieved if one could single out predictions that have a relatively low level of confidence and just ask users to make those decisions manually. This observation opens the door to the evaluation of more interactive scenarios and the exploration of tradeoffs between accuracy and the number of decisions where we might want to query the user – in other words tradeoffs between accuracy and user burden. While it is unrealistic to expect users to want to manually specify decisions on over 60 permissions (average of over 20 apps per user and over three permissions per app), it is not unreasonable to think that users might be willing to enter five to 10 decisions. In theory, if users were ready to enter all 60 decisions manually, one could theoretically reach 100% accuracy. The question is, how much accuracy do we lose by requesting users only to provide a fraction of these decisions.

Results presented in this subsection were obtained using the LibLinear tool for large-scale classification mentioned in the previous section. We use L2 loss logistic regression from LibLinear and compute labeling confidence measures for each test data point. The classifier provides the same accuracy as that reported for FS-8 in Figure 4 (87.8%) while also estimating the probability of each class label. Accordingly, we can compute the confidence of a given labeling decision as $Confidence = |Prob(Label = +1) - Prob(Label = -1)|$ where the predicted label is the one that has the higher probability, either +1 ("Allow") or -1 ("Deny"). When the classifier provides a confidence score for a label that is below a certain threshold, the system can prompt to ask the user for a decision. If we set a low threshold, the classifier will have a higher level of accuracy overall and at the same time result in higher user burden.

Results obtained by varying the threshold level and adjusting the percentage of decisions (or "data points") where the user is queried (horizontal axis) are presented in Figure 3.3. Again, these results are obtained using 10-fold cross-validation. Figure 3.3 plots precision on "unlabeled data," namely on those decisions where we do not query the user, as well as overall precision, namely combining both predictions made by the classifier when confidence is above the threshold and predictions made by the user when confidence is below the threshold. We assume that, by definition, querying the user has 100% accuracy. As can be seen, when asking users to make just 10% of the permission decisions, overall accuracy climbs from 87.8% to 91.8%. Given that users have already installed four applications out of an average of about 20 and that an app requires an average of three permissions, this amounts to asking users to provide five permission decisions (10% of 48 app-permission pairs). If users were willing to answer 10 permission decisions, overall accuracy would jump to over 94%.

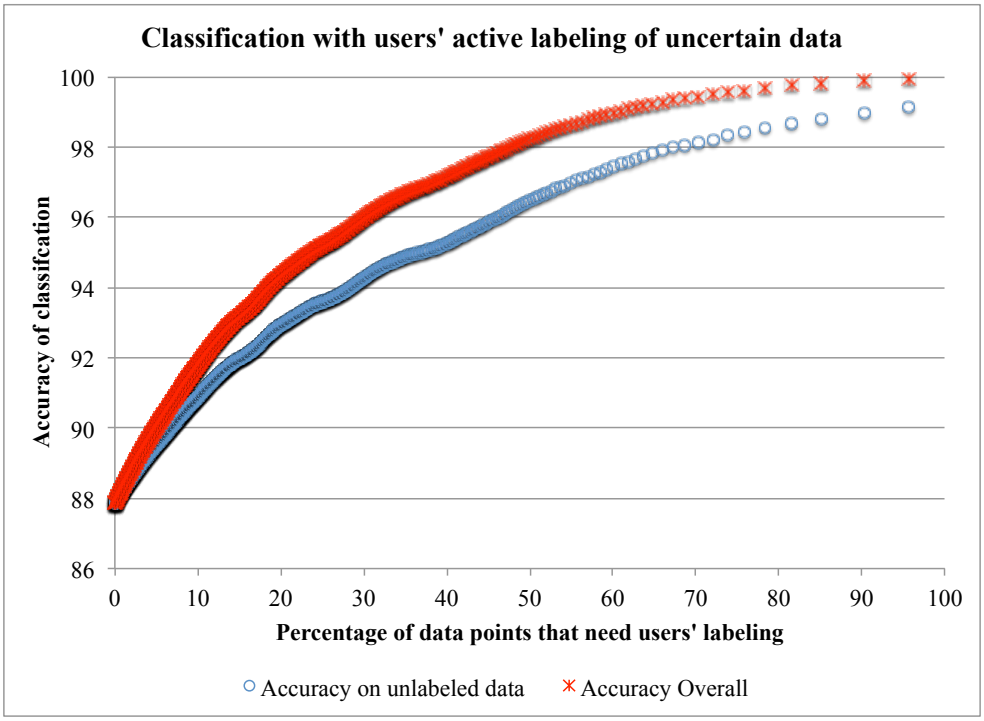

Figure 3.3: Accuracy Improvements by Interactions with the Users

3.4 Simplifying Privacy Decisions Using Privacy Profiles

As discussed in Chapter 2, previous work has shown that while users' privacy preferences are diverse, a relatively small number of privacy profiles can be identified, which collectively do a good job at capturing these privacy preferences. Each profile effectively corresponds to a different group or cluster of like-minded users and captures their privacy preferences. By asking users a few questions or presenting them with easy-to-understand descriptions of available profiles, it is possible to match individual users with profiles. In turn, these profiles can help predict with a high level of accuracy many of the users' location privacy preferences. A major motivation for our study of the LBE dataset is to determine to what extent mobile app privacy preferences, as captured in this dataset, exhibit similar patterns, namely to what extent a relatively small number of privacy profiles could be identified to simplify app permission decisions.

3.4.1 Generating Privacy Profiles by Clustering Like-Minded Users

Each user can be modeled as a vector of app-permission decisions. Such vectors are very sparse and did not yield the best predictive performance in our tests (see Table 3.1). Instead, aggregation of user preferences along each of the 12 permissions in the LBE dataset was shown to yield higher performance (FS-8). Accordingly, we represent each user as a 12-dimensional vector of

their aggregated preferences on each permission. Using a K-means algorithm with Euclidean distance, we proceed to identify clusters of users. This is done using the standard "cluster" toolbox in R for our implementation.

We now turn our attention to determining a good value of K, namely the number of clusters or privacy profiles. When comparing the different values of K, we consider three distinct metrics:

- **The accuracy of predicting default settings for users**
  As stated earlier, an important objective of our work is to determine to what extent a small collection of profiles can collectively help achieve a high level of accuracy.

  By simulating a scenario using as input user decisions on 20% of app permissions, we seek to discover how accurately the learned classifier is able to predict the rest of the users' permission settings.

  To this end, we rerun the classification task while replacing the identities of users with their cluster membership. The resulting loss in accuracy will tell us to what extent the profiles are collectively capturing the complexity and diversity of privacy preferences of our user population. We use the same 10-fold cross-validation procedure discussed in Section 3.3.2. We denote the average accuracy as $Accu(K)$.

- **Interpretability**
  This is a more subjective metric. Here, as we vary the number of clusters (K) we want to know to what extent we can still identify a small number of features that can be used to characterize each cluster. The idea is that these compact descriptions could possibly be presented to users who would then identify which profile best matches their preferences – based on a relatively small (and hence understandable) number of features.

  Here we define a quantifiable interpretability score in the process below:

  - We define $S(u, p, d)$ as the number of d decisions made by user u on apps requesting permission p. We define $S(u, p) = S(u, p, "allow") + S(u, p, "deny")$.
  - We define the users' agreement score $A(C, p, d)$ of all users in profile C making decision d (allow or deny) on apps requesting permission p:

$$A(C, p, d) = \frac{\sum_{u \in C} S(u, p, d)/S(u, p)}{\sum_{u \in C} 1}$$

  - We define the discriminative score $D(C, p)$ of permission p in profile C as:

$$D(C, p) = \max_{d} \frac{\sum_{C' \neq C} (A(C, p, d) - A(C', p, d))}{K - 1}$$

  For example, if we have three privacy profiles, and 99% of users in one of the profiles agree to deny access to the phone's location (across all apps), while 5% and 3% of the users in the other two profiles respectively agree to deny it, then we claim that the "Denying access to location permission" has a discriminative score of 95%.

- Then we define the interpretability of the choice of K $Interp(K)$ as:

$$Interp(K) = \frac{\sum_C \max_p D(C,p)}{\sum_C 1}$$

In short, highly discriminative features contribute to the interpretability of clusters. As indicated earlier, the thinking is that clusters that are easy to interpret would also make it easier for users to identify which cluster best matches their preferences. An alternative scenario might involve asking users discriminative questions to identify clusters that best match their preferences.

- **Stability of the profiles**
  Stability is yet another desirable attribute of clusters. We do not want our privacy profiles to change in response to a small perturbation in the data. We compute a stability metric based on the following algorithm. Given a collection of privacy profiles obtained for a given value of K, we randomly split all the users into 10 groups of equal size. We then use each possible combination of nine groups (of users) as training data for our K-means algorithm. For each combination of nine groups, we use the resulting cluster centers to relabel all the users.

  This gives us two sets of cluster labels for the same group of users: the original labels and the ones obtained from the relabeling. We use maximum-weight matching of bipartite graphs to find the mapping between the two sets of clusters. A stability score can then be computed as the percentage of users who remain in the same cluster. The stability score of the original privacy profiles obtained for a given value of K, denoted $Stab(K)$, can be defined as the average stability score taken across all combinations of nine groups.

These three factors mentioned above can help us choose a proper number of clusters, which is usually subjective to application scenarios. Here we propose a sample way to compute an overall score for each value of K:

- We assume that if a user's profile assignment is not stable, the system would simply use a classifier that does not take profile assignment information as input. We define $AvgAccu$ as the accuracy of such non-personalized classifier. Then we define the adjusted accuracy $Accu'(K)$ of the profile-based solution as:

$$Accu'(K) = Accu(K) \cdot Stab(K) + AvgAccu \cdot (1 - Stab(K))$$

- We define the overall score for each value of K:

$$Score(K) = \frac{2 \cdot Accu'(K) \cdot Interp(K)}{Accu'(K) + Interp(K)}$$

While imperfect, this metric can help us compare the benefits associated with different numbers of clusters (namely different values of K). Results obtained using the LBE dataset, including all four metrics, are shown in Figure 3.4. From the results in Figure 3.4, we can see that:

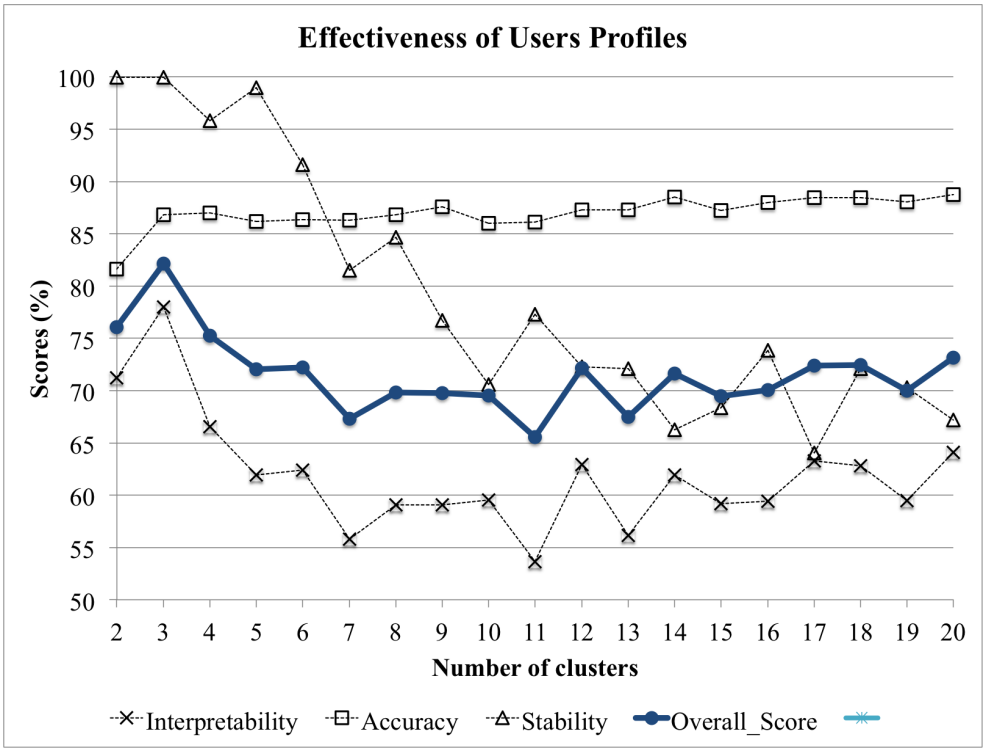

Figure 3.4: Effectiveness of User Profiles: Scores Under Different Ks

- As the number of clusters increases, classification accuracy gradually increases too, especially for $K \geq 3$. These clusters with increased prediction power can better sort out the differences between users in terms of their preferences. The accuracy of this form of lightweight personalization theoretically converges to that of the fully personalized method captured with the feature set FS-10 in Section 3.3.2. Results shown in Figure 3.4 indicate that somewhat similar performance can be achieved with values of K as low as 3.
- As one would expect, the stability of our clusters moves in the opposite direction, decreasing as the number of clusters increases.
- The interpretability scores of privacy profiles fluctuate as K changes, with a rapid drop beyond K=3, as the clusters become finer and more difficult to articulate. At the same time, we believe that this metric should be taken with a grain of salt. User studies would be needed to evaluate this issue better. In addition, it is likely that a simple wizard could easily be built to sort people among a set of available clusters/profiles by asking them a small number of questions. A user who has already installed a number of applications and configured permissions for these applications could be classified as falling in a given cluster without even having to answer a single question.

Though K=3 shows an optimal score under our quantitative definition, we are inclined to believe that the interpretability metric used above is somewhat simplistic and that usable solutions could

be developed for somewhat higher values of K, which in turn could yield higher accuracy levels.

3.4.2 Capturing the Discriminative Features of Each Cluster

With the definition described in the section above, we generate discriminative scores of permissions for each profile of users. Figure 3.5 provides discriminative descriptions of profiles/clusters for K=3 and K=6. These discriminative features could provide the basis for asking users a few questions to determine which cluster they fall into.

Beyond the discriminative features depicted in Figure 3.5, it is also possible to visualize and compare different privacy profiles using different color schemes. Figure 3.6 shows such a representation for scenarios where K=3, K=4, K=5, and K=6. Each cluster is represented by a 12-dimensional vector, with each cell colored according to the cluster's propensity to allow or deny the corresponding permission. Dark blue denotes a strong propensity to grant the permission, dark red a strong propensity to deny, while white denotes a split population – or at least a population whose decisions range about evenly between "allow" and "deny" across all mobile apps. Judging solely from the color schemes, one would conclude that clusters for K=3, 4, and 5 are very distinct, whereas the value of adding a sixth cluster (K=6) is starting to become less prominent. All scenarios seem to have one cluster that is particularly conservative when it comes to granting permissions (C3 for K=3, C4 for K=4, C5 for K=5, and C3 for K=6). Starting with K=6, a second conservative cluster (C4) is starting to emerge, though its population is not quite as reticent as that in C3. In general, we see that some clusters of users appear somewhat lenient, while others are more conservative. As the number of clusters increases, the nuances become finer. Some permissions also seem to yield more diverging preferences than others. For instance, looking at Figure 3.5, it can be seen that "Positioning" elicits very different reactions in clusters C3/Profile3 and C4/Profile 4, for K=6.

Figure 3.7 shows the variances of user privacy preferences for each permission in each profile for different values of K. As expected, variance tends to decrease as the number of clusters or profiles increases. For K=1, namely a single one-size-fits-all profile, the average variance of all permissions is 0.511. In contrast, for K=5 (namely five profiles), the average variance drops l to 0.216.

As we can observe from the visualizations and the accuracy results, instead of undertaking fully personalized learning, we could potentially have simpler models in which the diversity of users' preferences could be visualized as above. The model could have a similar level of prediction power (86.8% compared to 87.8% on prediction accuracy) using a small number of profiles.

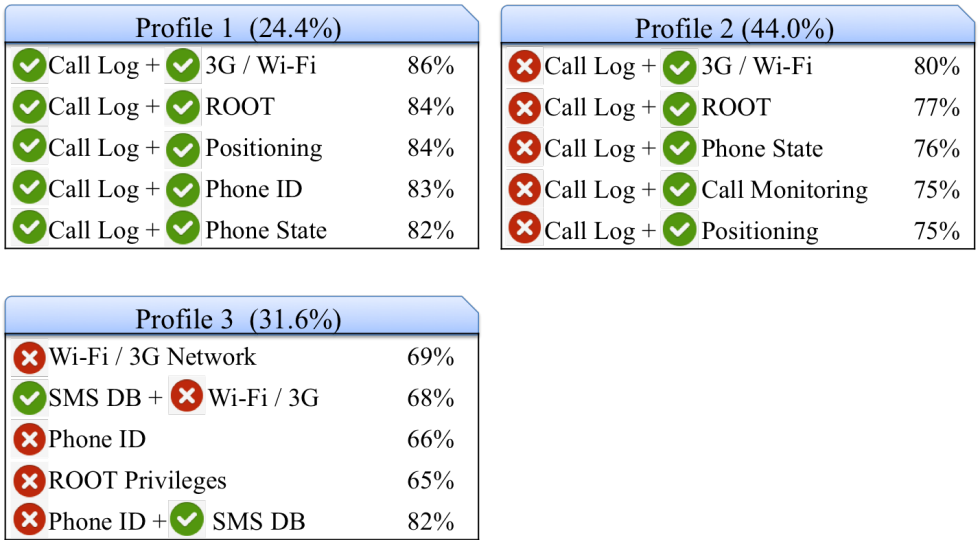

(a) K=3

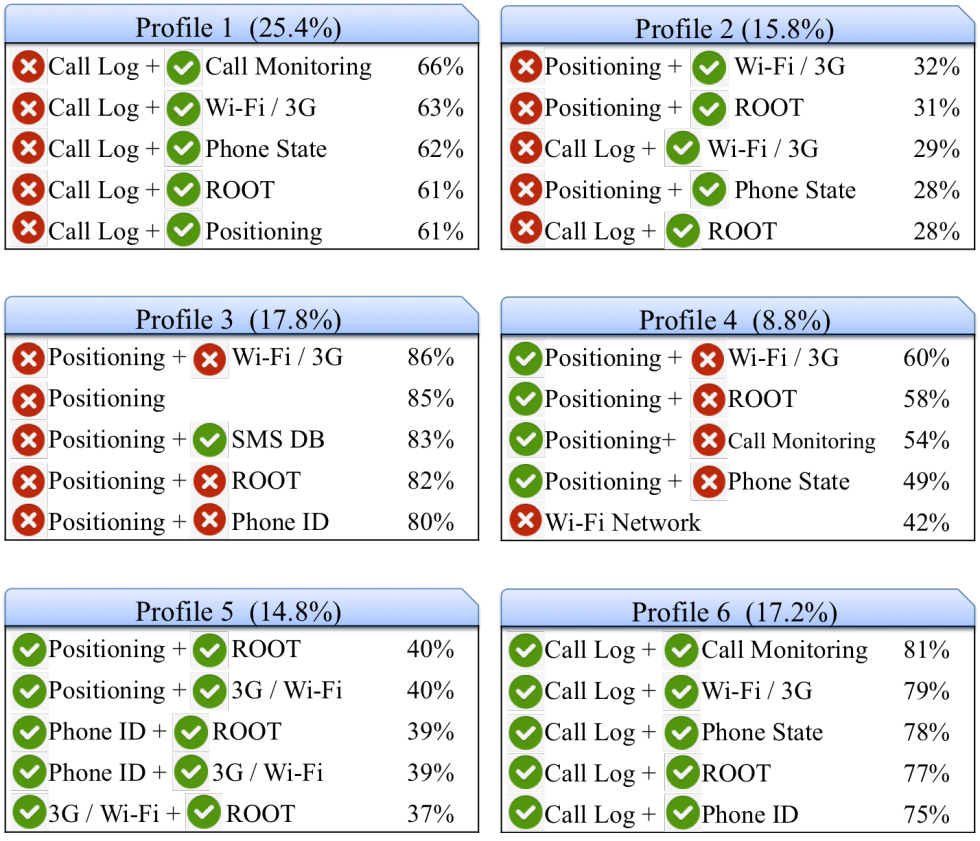

(b) K=6

Figure 3.5: Discriminative Descriptions of Privacy Profiles

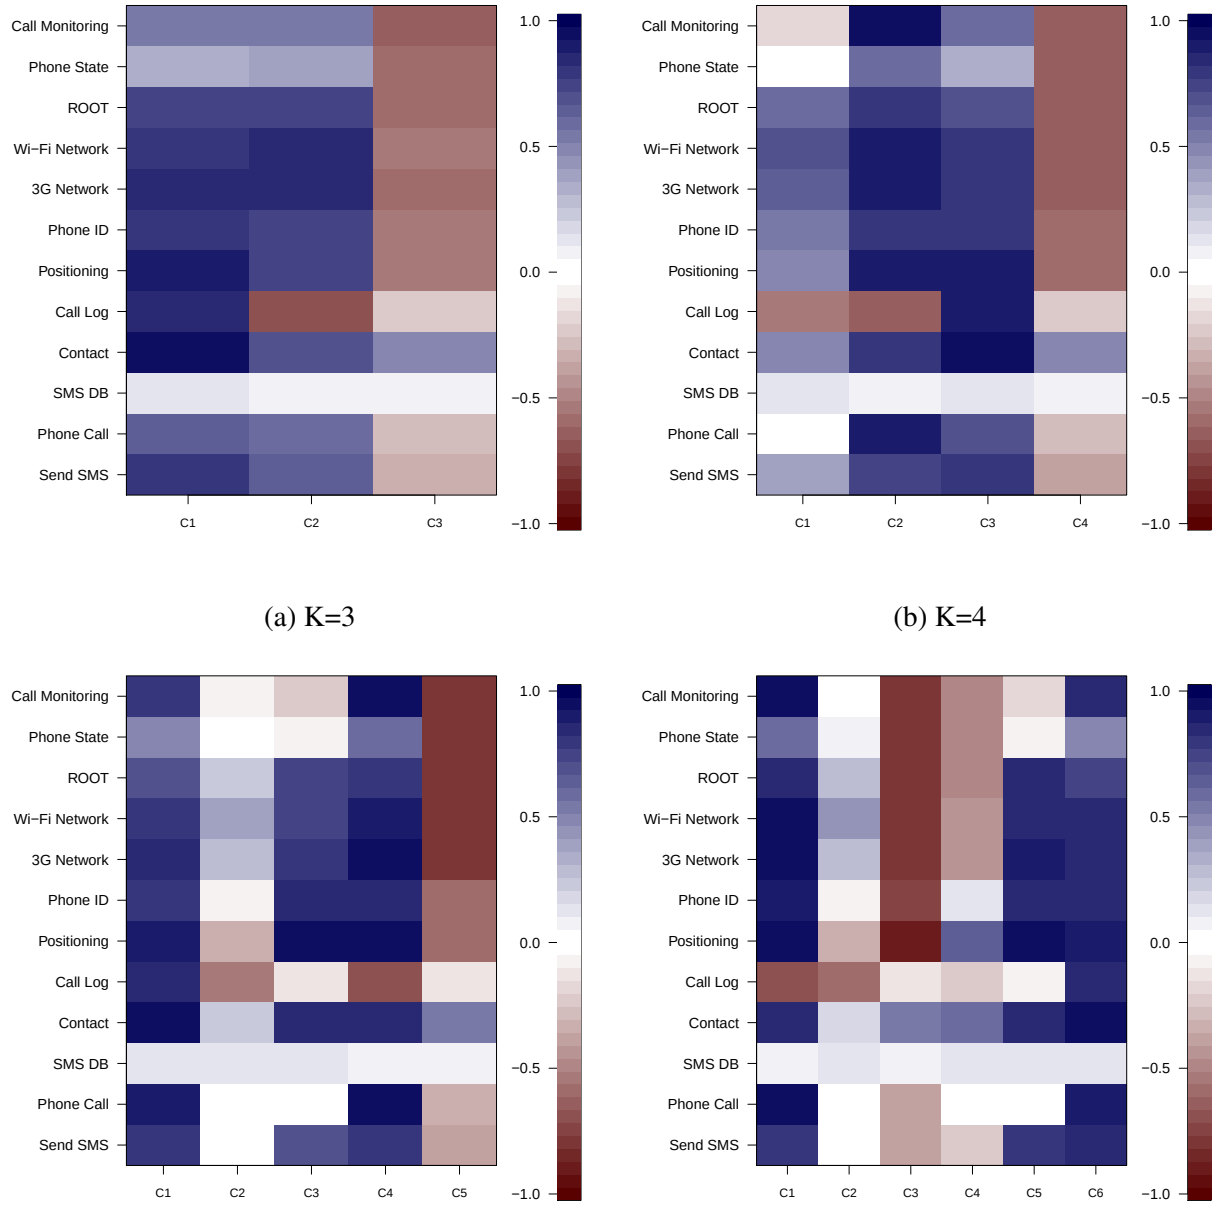

(a) K=3

(b) K=4

(c) K=5

(d) K=6

Figure 3.6: Colored Heatmap of Average Preferences in Each Privacy Profile. The color represents the average preferences of users in the corresponding cluster (horizontal axis) on the permission (vertical axis). For example, if a cell in the matrix has a value close to "-1" (mostly deny), then most of the users in the cluster can be expected to deny access to the corresponding permission requests.

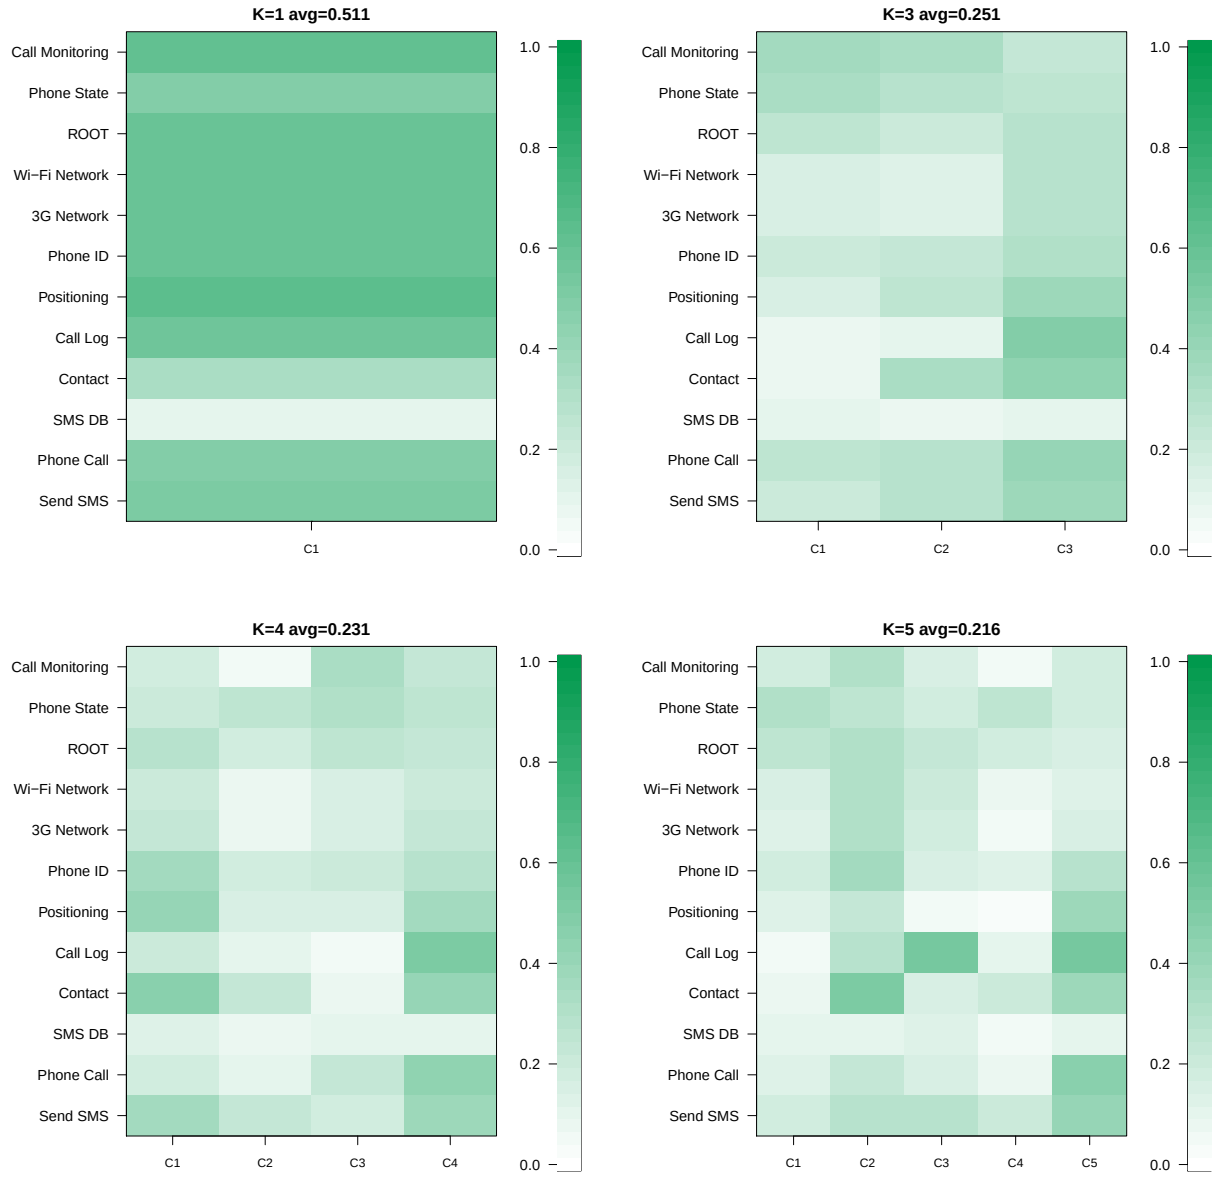

Figure 3.7: Variances of Preferences in Each Privacy Profile (see Figure 3.6). The darker the color, the higher the variance. "K=1" represents the case if no clustering is performed.

3.5 Discussion

In this study, we were the first to look specifically at actual settings of smartphone users on a large scale. We collected data from 4.8 million users and analyzed 237 thousand users who actively engaged in managing their permission settings. At the same time, we acknowledge that the data of their permission settings may not perfectly reflect the preferences of the general Android user population:

- The LBE Privacy Guard app is a popular app on various Android vendors, with 4.8 million users. Most of the users were expected to have this app pre-installed on their MIUI phones. Thus, users were not explicitly notified or educated about how to use this tool. We did a preprocessing of the dataset to retain only the permission settings contributed from active users on relatively popular apps on the Google Play store. However, even under this circumstance, we still observed that a fraction of the users had relatively permissive preferences, which is consistent with Westin's segmentation[55].
- App behavior changes across different versions of Android. The LBE app users were early adopters of a mobile app permission manager before Android introduced App Ops (Android 4.3, 4.4.2) and later "App Permissions" in the settings (Android 6+). Apps in their early versions might not handle the denial of permission settings properly to maintain the usability of the app without the requested permissions. Thus, we would expect that users may have different behavior when dealing with app permission in new or future Android environments. People's privacy preferences may also be slightly different across countries or regions[58]. Even though we have this limitation, we still found that a lot of LBE users had been actively managing their app privacy settings as early adopters and provided us with insightful inputs.

In the analysis, we showed that although the users' app privacy preferences could be diverse, it is possible to build models that could potentially predict users' tentative decisions for incoming permission requests from the new apps with fairly good prediction power. This opens the gate for designing and building tools that can help reduce user burden in managing the app privacy settings:

- As demonstrated in this chapter, the prediction model of users' app permission settings should be a personalized one instead of a one-size-fits-all solution. In the analysis, we simulated the scenarios by predicting users' app permission settings by learning from 20% of permissions settings from each user. Collecting sufficient input or feedback from the users before showing personalized recommendations is a common strategy for crowd-powered recommender systems. However, considering that privacy is a secondary task for users and relates to sensitive risks of unwanted sharing of private data, it is crucial to design a way to better capture users' diverse preferences effectively and efficiently before providing personalized decision support for users on app permission settings.

- These assistant tools can be further strengthened by enabling interactive learning with the users. In the analysis, we simulated this interactive learning process by first estimating the probability or confidence of the predictions for a certain permission request. In cases where the prediction model does not have sufficient confidence in the prediction, the tool selectively prompts to ask the user. By asking about only an additional 10% of the decisions, the accuracy of the prediction could climb to 92% from 87%. Traditional recommenders, such as shopping ads, usually interact with users in a passive way: they collect users' implicit feedback from clicks and annotations. However, in the scenario of assisting users on configuring their app permission settings, it is possible and convenient to interact with the users actively to better exploit the users' preferences.

From the dataset, we analyzed the diversity of users' preferences and performed simulated experiments to explore the potential of building a tool to help users with app permission settings management in order to reduce user burden while maintaining their privacy. The dataset is obtained from the anonymous tracking logs of a permission manager. A user-oriented study that captures users' reactions to such assistant tools could better help researchers to understand users' needs regarding permission management and iterate the design of such tools to better reconcile the privacy and usability of app permission settings management.

3.6 Summary

In this chapter, we obtained and analyzed a dataset of real Android users' app permission settings. From the results of the analysis, we showed that it is possible to significantly reduce user burden while allowing users to better manage their mobile app permission settings. In particular, we showed that personalized classifiers are able to predict users' app permission decisions. Considering the scenario where a user first installs and configures a small number of apps, we showed that using permission decisions made by the user for these apps, along with the app permission decisions from a representative population of users, it is possible to predict other permission decisions with an accuracy of over 87%. We further showed that by selectively prompting to ask users to make decisions on permission settings where the confidence of the prediction is below a certain threshold, the accuracy of prediction could climb to above 92% by asking about only an additional 10% of the decisions.

Also, our experiment of segmenting users into clusters of like-minded users shows that it is possible to describe users' diverse privacy preferences using a relatively small number of privacy profiles, while maintaining a similar level of prediction power compared with fully personalized classifiers. This opens the door to designing simpler mobile app privacy control interfaces, where users do not need to give up control in return for usability. Instead of waiting for users' sufficient inputs of app permission decisions, it is possible to first use a relatively small number of interactions with users to identify the privacy profile that tends to be closest to the users' preferences,

then provide personalized app permission settings.

Chapter 4

Using Privacy Nudges to Collect App Permission Settings of Engaged Users

4.1 Introduction

In Chapter 3, we obtained and analyzed a large-scale dataset of users' permission settings. This dataset helped us understand users' diverse preferences and enabled modeling and predicting users' app permission settings. However, one limitation of the dataset is that the data were collected in a passive way. Users were not educated or notified about the existence of the permission managing tools. As a result, a vast majority of the users of LBE Privacy Guard app were not actively engaging in configuring these settings. Would similar results be observed on data obtained from users who are motivated to think more carefully about their permission settings?

In this chapter, we explore a different approach that use nudges to motivate participants to actively engage with the app permission settings. Nudging has been shown to be effective at keeping users engaged in managing their privacy settings[13, 14]. Considering that we did not have the luxury to recruit a significant number of participants to use a permission manager over a long period of time, we used daily nudges to increase participants' awareness and motivate them to review and adjust the settings.

We conducted a field study with real Android users and collected participants' mobile app permission settings that they applied to their own devices in their daily lives. Specifically, we made the following design choices for the field study:

- We wanted to collect app permission settings that were actually applied to participants' phones. Previous studies (such as[59, 92]) use online surveys to ask about participants' comfort in allowing permission access to an app, or use text and images to simulate scenarios in which participants can specify their tentative decisions on allowing or denying. However, as Acquisti et al.[9] suggested, users might not configure their privacy settings

to be restrictive even if they claim to have high privacy concerns. Considering that none of these decisions from surveys were really applied to the participants' phones in their daily lives, the decisions could be somewhat aspirational and might not match their actual behavior[70].

- Our study app was designed to be more informative and easy to access. We developed an enhanced version of Android permission manager so that participants could use this tool directly to manage their actual app permission settings on their phones. At the time when we conducted the study, Android provided users with an app permission management tool named "App Ops." However, the tool is hidden by default. Users would need to install related utility apps to launch the screen for configuring permission settings. Our app used reflection hacking to utilize the same APIs called by App Ops. Instead of having low visibility, our app provides direct access to review and adjust the app permission settings. The app shows richer information to users for each permission request. It also displays frequency of access for each permission used by each app. In addition, we provided purpose information[57, 59] so that the participants could potentially configure the settings such that they would be comfortable with not only the permissions themselves but also the purposes associated with those permissions.
- We used daily privacy nudges to motivate the participants to engage with the permission settings. Considering that we did not have the luxury to recruit a significantly large number of participants to use our permission manager app over a long period of time, it was crucial for us to figure out a way to increase users' awareness of and engagement with management of their permission settings. We introduced a methodology that relies on nudging users by showing them pop-up messages that inform them about the information collected and shared by their apps. In this study, we aimed to reach a state where users' privacy settings are relatively stable. We did this by sending daily privacy nudges to users so that they were able to be aware of permission accesses and review or modify their settings over the course of a week. By using the nudges, we achieved greater certainty that the settings collected from participants were from users who had been motivated to engage with them.

We collected Android permission settings from 84 participants during a two-week study using the enhanced permission manager app. We recruited participants using rooted Android phones so that the permission manager was able to change the settings for the user directly onto their phone during the study. Considering the limitation of lab resources, it was impractical for us to collect a dataset at a scale similar to the LBE dataset. In this chapter we present a down-sampling analysis of our methods on a large-scale dataset and show that this technique is able to support the development of privacy preference profiles with strong predictive power (90% F-1 score) based on a fairly small sample of users.

In this chapter, we report the design of our permission manager app, the methodology we used to build clusters of like-minded users, and the findings from the analysis of the learned

privacy profiles from the clusters of users.

4.2 Data Collection Using Enhanced Android Permission Manager

4.2.1 Permission Manager App Design

Starting from the earliest versions, Android has used a permission-based mechanism to organize phone resources and APIs. With more advanced sensors and functionality integrated into the phones and more advanced APIs supported by newer versions of Android, the number of permissions has been increasing. By Android 6.0, the number of permissions had reached 235. LBE Privacy Guard was one of the earliest apps that implemented permission control by intercepting corresponding system calls of permission data.

With the introduction of version 4.3, Android officially provided the functionality of a permission managing tool by introducing a hidden app named "App Ops" (see Figures 4.1 and 4.2). This app was not accessible to users unless they installed a utility app that could bring the App Ops UI to the foreground of the phone. The Android system uses a class named "android.app.AppOpsManager" to manage the individual permission access. Each permission request from an app can have four modes: the system grants the data access if the mode is MODE_ALLOWED, ignores the request if it is MODE_IGNORED, ignores the request and raises an exception if it is MODE_ERRORED, and does nothing if it is MODE_DEFAULT. Only MODE_ALLOWED (allow) and MODE_IGNORED (deny) are commonly used in practice. The "android.app.AppOpsManager" requires system privilege (aka root access) to effectively edit permission settings.

In this field study, we developed our own Android permission manager app to collect users' app permission settings that may reflect their app privacy preferences. The participants we recruited for the study were mostly using version 4.3 and 4.4.2 Android phones. We implemented the serving logic of reading and modifying app permission settings by making API calls to "AppOpsManager."

To increase users' awareness and engagement, so that they review their permission settings if they find a setting they do not agree with, we made a number of modifications and enhancements to the Android permission manager App Ops, which we describe below:

- **We provided users with easy-to-access permission control.** Specifically:
  - In the stock Android permission manager, users had to configure individual permission items one by one (as shown in Figure 4.2(a)). We simplified the design of permissions in the enhanced permission manager app used by the field study. It or-

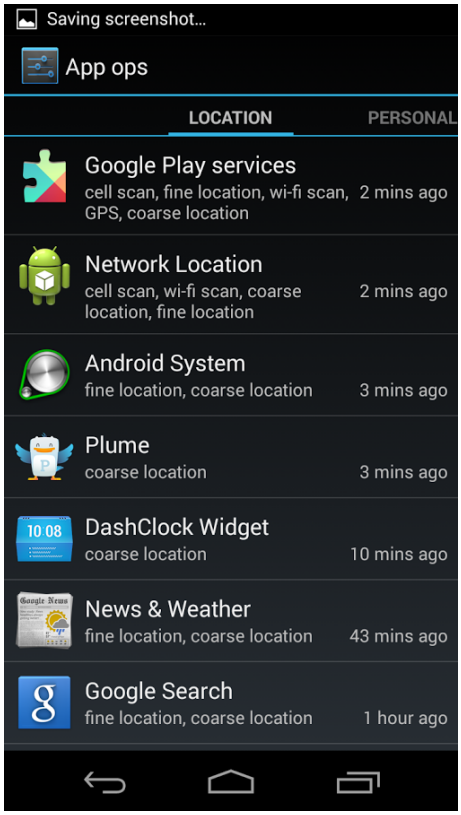

(a) The App Ops home screen. There are in total four groups of permissions managed in this app: Location, Personal (such as contacts and accounts), Messaging, and Device (such as camera). Each page shows the list of apps accessing a certain type of permission. If the user clicks a specific app in the list, App Ops will show a screen of detailed control options for that app (see Figure 4.2).

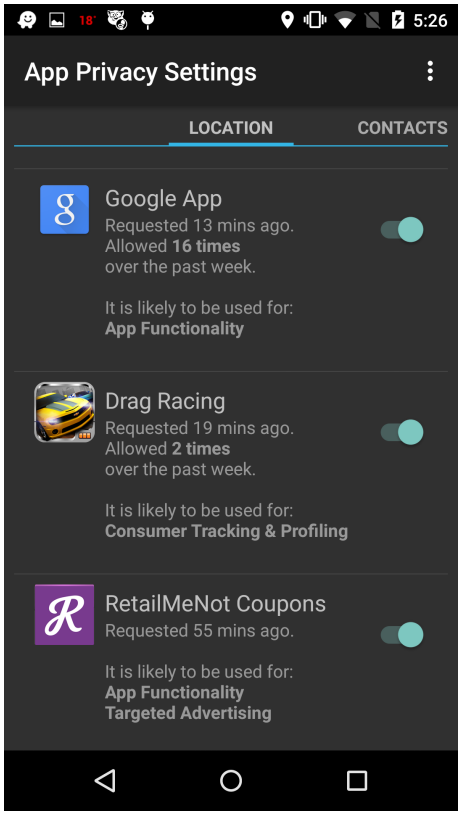

(b) The home screen of our enhanced permission manager. We provided six different groups of permissions for users to configure: Location, Contacts, Messages, Phone (phone call logs), Camera, and Calendar. Different from App Ops, users could easily configure their settings directly on the list screen. In addition, the app showed information such as the frequency of access (monitored by the permission manager) and the potential purposes of permission use[59].

Figure 4.1: The App Ops User Interface (Source: https://play.google.com/store/apps/details?id=fr.slvn.appops and the enhanced permission manager used in our study.)

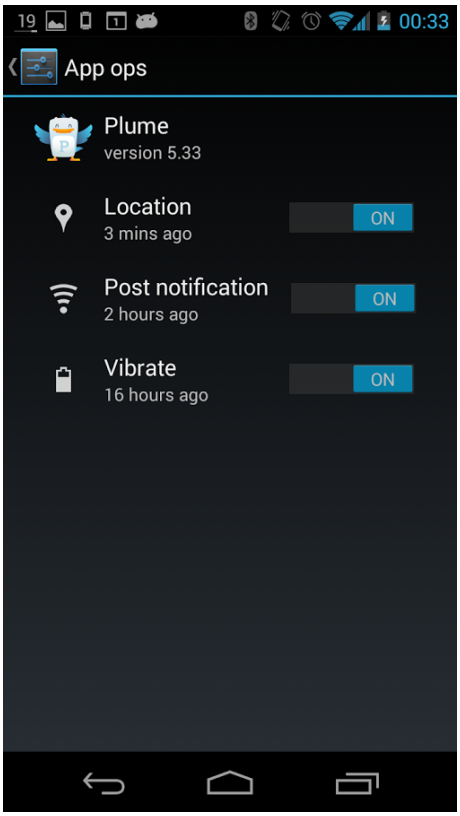

(a) The detailed control screen if a user clicks a specific app in Figure 4.1(a).

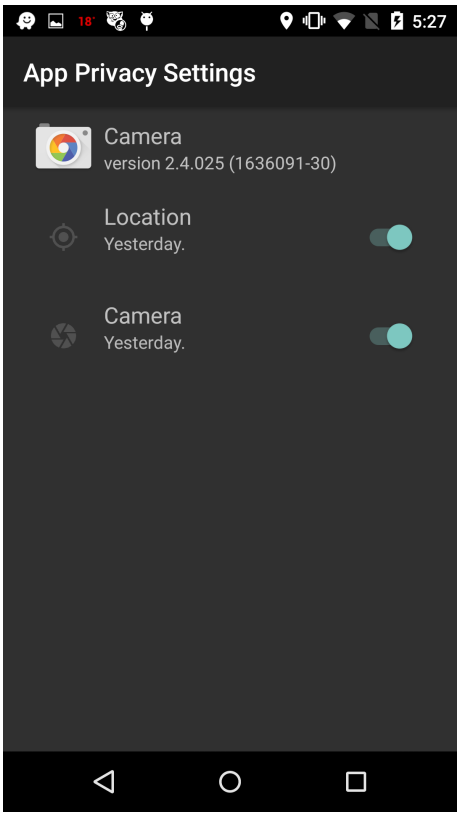

(b) The detailed control screen if a user clicks a specific app in Figure 4.1(b).

Figure 4.2: The App Ops User Interface (Source: https://play.google.com/store/apps/details?id=fr.slvn.appops and the enhanced permission manager used in our study)

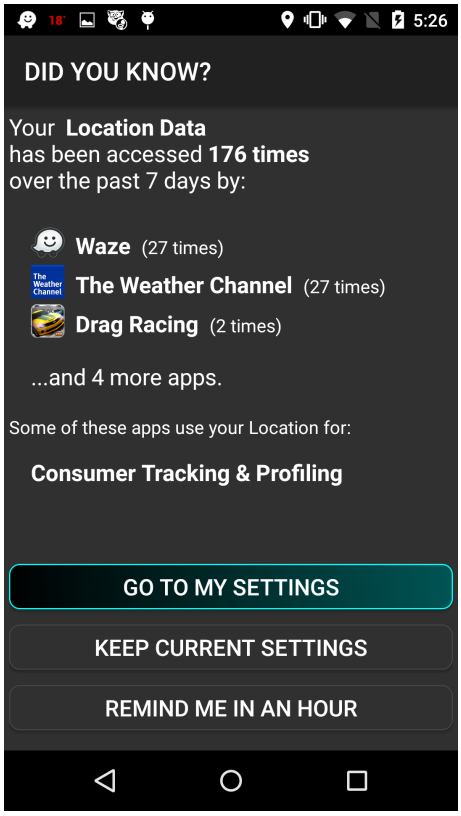

Figure 4.3: Daily Privacy Nudge Screen of Enhanced Permission Manager App. During the two-week collection of users' permission settings, starting from the second week when we had collected a week's worth of access frequencies of permission requests, we showed a daily privacy nudge every day between 12 pm and 8 pm. Each day, the nudge message would randomly choose a permission type and show the information about the apps that access this permission. Users were able to react to the nudge by going directly to the permission settings page, ignoring the message, or postponing the response.

ganized the permission items into six common privacy-related permission groups: Location, Contacts, Messaging, Call Log, Camera, and Calendar. As a result, multiple permissions were represented as a single permission in the study app, thereby giving users fewer setting items to manage. For example, READ_CONTACTS and WRITE_CONTACTS are represented as "Contacts." This grouping is partially based on results by Lin et al.[57] and Felt et al.[37]. Users can directly allow or deny each permission while reviewing it in the permission manager. Coincidentally, Google announced similarly grouped permissions for Android 6.0 shortly after we conducted the study.

  - The stock Android permission manager showed permission list and permission controls in separate screens (see Figure 4.1(a) and Figure 4.2(a)). Users would have to click twice and go back and forth to review and adjust the permission settings. In the study app, we put the control toggle button and the list of apps requesting a permission on the same screen (see Figure 4.1(b)). Users could choose to browse permission settings organized by permissions or by apps.

- **The permission manager used in the study showed richer information about each permission request.** The manager showed not only the most recent time when an app had accessed a permission, but also the frequency of access over the past week and the tentative purpose for collecting this type of data (see Figure 4.1(b)).

  Purpose information is not yet available for either Android or iOS ecosystems. We obtained the purposes for an app requesting permissions through static analysis of Android apk files. Lin et al.[57, 59] used Androguard to identify third-party libraries included in the app and scanned the permission usage of each library to infer likely purpose(s). An app can be labeled as using a permission for app functionality, targeted advertising, consumer tracking and profiling, and/or sharing with social network services. This information about purpose can enhance awareness on the part of users. However, users would still need to control the permission access on a per-app basis, instead of a per-purpose basis (such as granting access to an app for the purpose of delivering its core functionality, but restricting the ad libraries in the app from getting this data for the purpose of advertising).

- **We used daily privacy nudges to motivate users to engage with the permission manager and review/revise their permission settings.** As we discussed at the beginning of the chapter, we found that it is important to motivate users to engage in configuring their permission settings and keep them aligned with their preferences. Privacy nudges have been found to be effective towards this goal[14, 26]. We adopted a similar strategy to show one privacy nudge per day to each study participant (see Figure 4.3). The detailed information shown in each nudge is about a specific permission. Each nudge displays access frequency for each app and likely purpose(s) of apps accessing permission data. Participants were able to directly open the permission manager to adjust the settings, close the nudge and keep the current settings, or postpone the nudge message and have it be shown

again in an hour.

It is also important to note that a couple of months after we finished the field study in this chapter, Android 6.0 released a new design of Android permission manager. Starting from Android 6.0, a runtime permission request dialog has been available with backend permission management logic implemented in "`android.content.pm.PackageManager`." The PackageManager class also maintains flags recorded for each permission setting, such as whether the decision has been made by the user or device policy, whether the permission setting is permanent or needs to be prompted again, and so on.

4.2.2 Study Protocol

In this section, we explain the protocol we followed when conducting the study to collect users' app permission settings and other feedback.

Since permission management requires system-level privileges, this study had to be conducted with users of rooted Android phones. Importantly, our participants installed our app on their own rooted Android phones – namely the phones they use in their regular daily activities. While users of rooted Android phones may constitute a biased population, this approach still allows us to evaluate the practicality of building privacy settings profiles and using a Personalized Privacy Assistant on real users. Assuming it will be possible to customize permission management in future versions of mobile platforms, the same approach can be adopted to build privacy profiles representative of the general population's privacy settings.

We recruited Android phone users who used a rooted Android phone (4.4.X or 5.X; Android 6.X had not been released at the time of the study) with a data plan. We only recruited Android users who had used their phones for at least one month. Considering that our target population is limited to users of rooted Android phones, we recruited participants from multiple online communities related to Android in general or rooted Android in particular on Facebook Groups, Google+ communities, Reddit subreddits, and tech forums. Our study was approved by Carnegie Mellon University's Institutional Review Board. We disclosed that the study app collected and managed Android app privacy settings as it would have root access to participants' phones. All participants had to be 18 years or older. We asked participants to complete an initial screening survey to verify that they matched the above criteria and to collect demographic information. Participants who qualified were sent a download link for our permission manager and a username to activate it.

In the first week of the study, participants could use the permission manager to deny or allow permissions selectively. Our app also collected the frequencies of permission requests for installed apps, which were shown in the permission manager. In the second week, the participants received a privacy nudge once a day, between 12 pm and 8 pm. We waited one week before showing daily nudges to allow participants to familiarize themselves with the enhanced

permission manager and to ensure that the privacy nudge messages contained meaningful access frequencies based on the behavior of participants' installed apps.

The privacy nudges provided information about one of six permissions available in the enhanced permission manager. The selection of which nudge to show was randomized to counter the ordering effects. If a particular permission type had never been accessed by apps on the participant's device (access frequency would be zero), another permission type was selected to be shown in the nudge instead. In Figure 4.3, an example of a nudge showing Location access can be seen.

After participants completed the study, we asked them to fill out an exit survey online, consisting of the 10-item IUIPC scale on privacy concerns[66] and an 8-item scale on privacy-protective behavior[64]. They were afterwards compensated with a $15 gift card. We further invited all participants to an optional interview, in which we explored their reasons for restricting or allowing different permissions, their comfort level concerning their permission settings, the usability of the enhanced permission manager and privacy nudges, and the utility of adding frequencies and purpose information. Those who participated in the optional interview received an additional $10 gift card.

In total, we collected data and survey responses from 84 Android users and interviewed 10 of them. The 84 participants originated from North America (66; 62 U.S.), Europe (10), Asia (7), and South America (1). Given the target population of rooted phone users, we expected our study population to skew towards young, tech-savvy males. Our expectation turned out to be true. Indeed, the majority of our participants were male (78 male, 6 female) and 18–54 years old (median 23). Among them, 8 had a graduate degree, 22 a Bachelor's degree, and 5 an Associate's degree; 30 had attended some college, and 19 had a high school degree or lower. Most commonly reported occupations were student (35), computer engineer or IT professional (8), service (5), and unemployed (5).

Participants exhibited relatively high privacy concerns, scoring high on the IUIPC[66] scales for control (median 6.33, mode 6.33, min 2.33, max 7), awareness (median 6.67, mode 7, min 4, max 7), and collection (median 6, mode 7, min 1.25, max 7). We also compared the responses from the participants in our field studies in the year 2016 with the ones from the Pew survey[64] of the general population in 2015. The participants in our study took more measures to protect their online privacy compared with the general population, as shown in Table 4.1. This suggests that our participants' privacy settings may be more conservative than those of the general population.

In total, we obtained 4,197 permission settings from 84 participants, reflecting their allow and deny settings for the six permissions in the enhanced permission manager. We filtered the dataset to only analyze permission settings for apps available in the Google Play store. It is important to note that all participants used Android 4.4.X or 5.X phones, where app permissions were granted as "Allow" by default when an app is installed. Later Android versions prompt users to

Table 4.1: Privacy Protective Measures of Our Study Populations Compared With the General Population. Questions and general population results are based on a Pew survey[64]. The scale is selected from PIAL7 Q11 of the survey. The ordering of the questions is randomized for each participant. "While using the internet, have you ever done any of the following things?" (Multiple Choices: Yes / No / Doesn't apply to me / Don't know) Here we show the percentage of "Yes" among all participants who chose "Yes" or "No." The Pew survey data was collected in the year 2015, comparable to the data we collected from the field studies in early 2016.

Population	Pew Survey	Data Collection Study	PPA Field Study
Used a temporary username or email address	30.86%	90.00%	92.75%
Added a privacy-enhancing browser plugin (e.g., DoNotTrackMe, Privacy Badger)	11.11%	67.09%	57.35%
Gave inaccurate or misleading information about oneself	28.57%	83.75%	78.79%
Set browsers to disable or turn off cookies	44.16%	61.54%	63.24%
Used a service that allows to browse the Web anonymously (e.g., proxy, Tor, or VPN)	11.84%	81.01%	83.82%
Decided not to use a website because it asked for real name	29.49%	66.67%	54.84%
Used a public computer to browse anonymously	15.00%	49.35%	44.92%
Used a search engine that doesn't keep track of search history	22.39%	71.25%	63.64%

"allow" or "deny" permission requests, thus making this pre-processing unnecessary. Also, we analyzed only those permission settings for which the corresponding app had been launched in the foreground at least once during the study, or if users explicitly denied or allowed an app's permissions. After filtering, our dataset consisted of 3,559 individual permission settings for 729 distinct apps.

4.3 Permission Settings Data Analysis

Of the 3,559 permission settings, 2,888 were allowed (81.15%, mean: 34.38 per user), which is the default choice, and 671 (18.85%, mean: 7.99 per user) were denied by participants. Call Log requests were denied the most (41.33%), while Camera access was allowed the most (95.07%). Of the permissions participants changed explicitly, 7.58% were re-allows of permissions they had previously denied. In the interviews, we asked participants why they did not deny certain apps, in cases where they re-allowed or just never changed an app's permission. The main reason for re-allowing a permission request, as mentioned by two interviewees, was that denying it broke or might break app functionality. P6 noted, "The moment I turned it off I realized that it wasn't gonna send me any messages." Nine interviewees reported not denying permissions because they were required for the app to function. Two interviewees noted that they trusted the app or the app provider. P2 stated, "This fitness app is made by Google and I trust it so I allowed it."

We fitted the users' settings data to a random effect logistic regression model grouped on users' allow/deny decisions on app permissions by user ID. The independent variables include major features such as user demographics and app category, which could be obtained from our dataset. We retrieved the app category information from the Google Play store. The detailed logistic regression results are shown in Table 4.2.

App categories and the type of permissions can help predict individuals' allow/deny decisions, according to the p-values in Table 4.2. Participants mostly agreed on permission settings for certain app categories. For example, apps in the "Books & Reference" category were always denied access to Contacts and Call Log, while "Photography" apps were always allowed access to Camera, as is to be expected. In aggregate, participants' settings on app categories are somewhat diverse (average SD=0.388, if we define allow=0, deny=1). The detailed effect size (odds ratios) can be found in Table 4.2. Eight interviewees mentioned that they denied access based on app functionality, e.g., when the use of the permission was not clear or when they thought that an app would not need it. P4 stated: "I do not use Facebook for any calendar function, so I denied it access to my calendar." Four interviewees mentioned denying apps when they did not use them, especially pre-installed apps they did not uninstall. The permission type also played a role in how participants reacted to the nudges. They reviewed their settings mostly when the nudge was about Location access (25%), followed by Messages (23.75%), Call Log (18.75%), Camera (15.19%), Calendar (14.29%), and Contacts (12.20%).

The interviews provide insights into participants' reasons for denying apps permissions. Nine interviewees (out of 10) confirmed the usefulness of access frequency information; four stated it as a reason to deny a permission request, five mentioned it was useful in the nudge, and two stated it was useful in the permission manager. For example, P1 stated: "Didn't notice that the app had actually accessed the location that many times. It is pretty crazy." However, despite reported usefulness, we did not find a significant impact of access frequency on users' decision of permission settings (see Table 4.2).

Table 4.2: Random Effect Logistic Regression on Users' Propensity to Allow or Deny Permission Requests.

Factors		Odds Ratio	StdErr	z	$P > \|z\|$
Age		1.024816	.0619711	0.41	0.685
Gender		.6941319	.6480886	-0.39	0.696
Education	Associate	6.351436	6.536207	1.80	0.072
	Bachelor	.3252345	.2102106	-1.74	0.082
	Graduate	2.265247	2.258762	0.82	0.412
	High School	.9914089	.5819914	-0.01	0.988
	No High School	1			
	Some College	1			
Occupation	Administrative	5.442226	8.371201	1.10	0.271
	Art/Writing/Journalism	1			
	Business/Management/Finance	1			
	Computer/IT	1.364362	1.553644	0.27	0.785
	Decline to Answer	5.775118	6.803399	1.49	0.137
	Education	.0920523	.1597209	-1.37	0.169
	Engineer in Other Fields	16.96705	31.93771	1.50	0.133
	Homemaker	1.134727	3.123314	0.05	0.963
	Legal	.1008037	.1688665	-1.37	0.171
	Medical	.633246	.8901533	-0.33	0.745
	Other	1.804592	2.601707	0.41	0.682
	Scientist	1.903118	2.983608	0.41	0.681
	Service	1.962722	2.268031	0.58	0.560
	Skilled Labor	.7758243	1.22502	-0.16	0.872
	Student	2.534309	2.248981	1.05	0.295
	Unemployed	1			
IUIPC Scale	Control	.6704036	.3212597	-0.83	0.404
	Awareness	.6779195	.381246	-0.69	0.489
	Collection	1.810677	.4923613	2.18	**0.029**
	Books & Reference	12.19531	9.009827	3.39	**0.001**

	Business	11.00032	6.011878	4.39	**0.000**
	Communication	4.464244	1.614809	4.14	**0.000**
	Education	5.988742	6.630343	1.62	0.106
	Entertainment	7.792989	3.563787	4.49	**0.000**
	Finance	3.490802	1.561327	2.80	**0.005**
	Game	8.974919	4.578022	4.30	**0.000**
	Health & Fitness	4.637063	2.497553	2.85	**0.004**
	Libraries & Demo	2.107152	2.378477	0.66	0.509
	Lifestyle	4.278822	1.932977	3.22	**0.001**
	Media & Video	5.627252	3.56555	2.73	**0.006**
	Medical	1			
	Music & Audio	14.15537	7.885298	4.76	**0.000**
	News & Magazines	6.177335	3.068304	3.67	**0.000**
	Personalization	.6819545	.5712842	-0.46	0.648
	Photography	1.099871	.8050647	0.13	0.897
	Productivity	2.107637	.8318742	1.89	0.059
	Shopping	4.381211	1.813481	3.57	**0.000**
	Social	7.208478	2.76813	5.14	**0.000**
	Sports	25.32193	17.04635	4.80	**0.000**
	Tools	3.562823	1.293064	3.50	**0.000**
	Transportation	.8090313	.530982	-0.32	0.747
	Travel & Local	1			
	Weather	1			
Permission	Location	2.620968	1.041181	2.43	**0.015**
	Contacts	.7826907	.3259032	-0.59	0.556
	Messages	3.870752	1.591046	3.29	**0.001**
	Call Log	2.39916	1.127688	1.86	0.063
	Camera	.1410928	.0698829	-3.95	**0.000**
	Calendar	1			
log(Frequency+1)		.9541353	.0317826	-1.41	0.159
Purpose	App Functionality	1.296318	.2925215	1.15	0.250
	Targeted Advertising	1.235337	.5431015	0.48	0.631
	Consumer Tracking & profiling	1.123383	.6212463	0.21	0.833
	Social Networking Services	.2956021	.3464561	-1.04	0.298
(Constant)		.0275754	.0780506	-1.27	0.205
Logged variance of random effect		.7827504	.2309066		
StdEv. of random effect		1.479013	.170757		
ρ (Intraclass correlation)		.3993685	.0553883		

Note: Likelihood ratio test of $\rho = 0$: $\bar{\chi}^2 = 338.10$, $P >= \bar{\chi}^2 : 0.000$.
We fitted the model with equal weight on all permission settings. The model also took input of participant IDs of each corresponding permission setting as "random" variable to counter subject-specific effects. The odds ratios (OR) show the associations between the specific factor and the user's decision: "$OR = 1$" means that they are totally independent; "$OR > 1$" means that with this factor taking effect, the user will be more likely to deny; "$OR < 1$" means that with this factor taking effect, the user will be more likely to allow. Due to the limit of the population size, some factors did not have enough data. The p-values ($P > |z|$), which show the significance of the corresponding association, are marked in bold if they are smaller than 0.05.

The logistic regression model indicates that purpose information by itself was not sufficient to predict whether a permission is denied by users in our dataset. This suggest that people's privacy preferences are generally more complex, and depend on more than just one factor such as purpose. Instead, if we are to accurately capture people's preferences and predict them, we will likely need models that look into combinations of factors (e.g., app category, permission, and purpose). For the sake of this particular study, we relied on purpose information obtained from the work of Lin et al.[59]. It should be noted that it does not provide purpose information for all apps and all purposes. Specifically, purpose information was available for 8.6% of apps requesting Location access, 35.1% for Contact, and 42.5% for Camera requests. Of the daily privacy nudges, 60.4% were shown with purpose information; 31.45% of those nudges showed purposes other than required for app functionality. Participants reviewed their settings as a reaction to 23.91% of the nudge dialogs mentioning "Targeted Advertising" and 17.77% of other types of nudge dialogs (chi-square=0.9804, df=1, p=0.3221, effect size(odds ratio)=0.6877). Participants were less likely to deny if at least some purpose(s) were shown (denying 13.53% of requests compared with denying 19.95%; Chi-square=10.1793, df=1, p=0.0021, effect size(odds ratio)=0.6784), which matches the results obtained by Tan et al.[88]. However, as mentioned above, purpose itself cannot be sufficient to predict users' decisions (all p-values of the purposes are high in Table 4.2). Nine interviewees mentioned that purpose information was useful in one or more ways: as a reason to deny (three), useful as part of the nudge (seven), and/or useful in the permission manager (three).

It is important to note that the participants saw the purpose information of a permission request in a bundle, if the corresponding app used it for multiple purposes (such as using it for both app functionality and ads at the same time). In comparison, Lin et al.[59] asked the participants' attitude on each purpose individually. From our study responses, three interviewees mentioned a tradeoff when applications had more than one purpose stated. They wanted the app's main functionality that needed a permission, but did not like that it was being used for other purposes. P3 stated, "Snapchat is a tradeoff. Although I'm not happy they access my contacts for tracking, I think I will allow them to access my contacts because of the function they provide." Participants' choices were typically permissive in such cases. This suggests that the actual permission settings offered to users, which do not differentiate between purposes, are not well aligned with their privacy preferences.

4.4 Profile-Based Method to Model and Predict Users' App Permission Settings

In Chapter 3, we showed that it is possible to use a relatively small number of profiles to capture users' diverse privacy preferences and use profile information to predict their app permission settings. With the results we obtain in the previous section, we see evidence that predictive models are likely to require models that account for app category, permission, and purpose information.

In this section, we use this information to cluster like-minded users and generate privacy profiles for each cluster. Then we use the profiles to predict permission settings of users in each cluster.

- Clustering of like-minded users: Given the dataset of users' app permission settings, we extract features that describe users' privacy preferences, such as aggregated preferences of users on specific app categories accessing permission for specific purposes. We group users who have similar app permission settings into a small number of clusters. For each cluster, we generate a privacy profile of predicted permission settings for all users in the cluster.
- Given a new user, we can first estimate the cluster that has the most similar preferences on app permission settings compared to this user. Then we can use the profile of this cluster to predict this user's app permission settings.

4.4.1 Clustering Like-Minded Users

From the dataset we collected in Section 4.3, we obtained users' app permission settings as a collection of rows in the form of $(user, app, permission, decision)$. Here we only analyze the apps available on the Google Play store. For each app, we obtained the app category information from the Google Play store page of the app. We obtained the purpose information on permission requests from the work of Lin et al.[59], which provides an indication of the purposes an app may use requested data for, but does not provide purpose information for all apps or permission requests. We quantify each user's preferences as a three-dimensional tensor of aggregated preferences of (app category, permission, purpose). For each cell, we define the value as the tendency of the user to allow or deny a given permission requested by apps from a specific category with a corresponding purpose: from -1 (100% deny) to 1 (100% allow), and N/A if we do not have the user's settings data for a cell. For each user u, app category c, permission p with a specific purpose r, if user u chose "Allow" a times and "Deny" d times for all apps within category c requesting permission p for purpose r, we define the user's tentative decision as:

$$P[u][c,p,r] = \begin{cases} \frac{a-d}{a+d}, & \text{if } a+d>0 \\ \text{N/A}, & \text{if } a+d=0 \end{cases}$$

In order to estimate similarities among participants' feature tensors, we applied tensor factorization to impute the missing values in the tensors. Compared with aggregating on specific dimensions, the imputation method does not bias towards any dimension mathematically. Also, to fit each imputed tensor so that it fits the original tensor on the observed values, we applied weighted PARAFAC tensor factorization[7]. We put 1-weight on all known data cells and 0-weight on unknown data cells in the tensor. Thus, we optimized the overall error of the imputed tensor in Frobenius norm regarding the observable data points from the dataset only. Next, we computed the numeric distances between the imputed tensors that represent the overall app permission preferences of users.

We chose hierarchical clustering as our clustering algorithm. In Chapter 3 we chose K-means because it worked well when we have a large number of users. Comparing to K-means and DBSCAN, hierarchical clustering is not sensitive to the size or density of the clusters. Thus, it allows clustering results with unbalanced sizes and can identify sub-groups in datasets of a smaller scale. We show the generated clusters in Figure 4.4. These clusters will be used in another field study in Chapter 5 to help users configure their permission settings, using what we will call "Personalized Privacy Assistant" functionality.

As already discussed, these clusters were generated from a dataset of permission settings collected from 84 rooted Android users who used the enhanced Android permission manager during our two-week field study. Obviously different parameter settings for clustering or predicting users' settings might lead to slightly different results. The parameters selected for the models displayed in this section are those that correspond to the highest five-fold cross-validated F-1 score on permission settings predictions (F-1 score=90.02%, hierarchical clustering: K=7, complete linkage, cosine distance). For reference, Figure 4.5 also shows results for different parameter combinations. From this figure, we can see that when including purpose information in the feature set, K=7 yields the highest F1 score.

Figure 4.4 is a visualization of the seven clusters of users we generated according to their permission settings. The colors of the cells indicate ,how likely users in the cluster are to grant (blue) or deny (red) different permissions to apps in different categories. Darker shades indicate stronger preferences (e.g. darker shades of blue indicate that users are particularly likely to grant the corresponding permission to apps in the corresponding category). Empty cells (white) indicate that we did not collect any setting data from users in this cluster about this permission and app category.

From the visualization in Figure 4.4 we can see that users in Cluster 1 are mostly permissive, whereas users in Cluster 2 are mostly protective when configuring permission settings. Obviously some entries have less supporting data than others. For example, the cell corresponding to medical apps accessing SMS data in Cluster 1 only has one entry, which happens to be a deny decision. One would be hard-pressed to use this one data point to suggest that all users in this particular cluster would want to deny this permission to all apps in this particular category. On

the other hand, some entries are supported by a large number of data points (e.g., over 200 data points from apps in the Tools category requesting Location permission in Cluster 1). Entries showing a very large majority of users lending towards granting or denying a permission are entries that are likely to lend themselves to the generation of recommendations for users in the same cluster (e.g., granting Location access to Tools apps for users in Cluster 1). Some entries have data that is less clear-cut. For instance, Tools apps requesting access to Location show that users in Cluster 3 are not always consistent when it comes to deciding whether or not to grant this particular permission to apps in this category. This, in turn, would suggest that it might not be safe to use this data to make a particular recommendation to users in this cluster when it comes to this particular permission and this particular app category. In an ideal world, one would want to collect more data from people in this cluster and build models that support making recommendations for people who fall in this cluster. Users in Cluster 3 to 7 seem to correspond to what Lin et al.[59] refer to as "Advanced Users" who exhibited more nuanced privacy concerns about specific app categories or permissions rather than having coarser preferences that extend across entire categories of apps or permissions.

Among the participants in our dataset, we found from the profile visualizations that many of them are assigned to permissive profiles. We observed that even though participants in this study were more tech-savvy than the average population, this did not seem to imply that they were more conservative when it comes to configuring their app permission settings. The apps installed on the participants' phones were selected by the participants themselves without our intervention. Accordingly, it is reasonable to assume that participants generally limited themselves to downloading apps to which they generally felt comfortable granting permissions. It is not entirely surprising that their selection of permissions would appear to be permissive if indeed they vetted their apps prior to downloading them.

Given the fact that the dataset we used for the clustering is collected from a user study of 84 users who have rooted Android phones, the profiles we generated may not reflect the general situation of the common Android user population. Nevertheless, our learned profiles do show a similar pattern on users' diverse preferences compared with the profiles or segmentation identified in previous works:

- Westin's index[55, 93, 94] categorized people generally into "privacy fundamentalists," "privacy unconcerned," and "privacy pragmatists." Among our learned profiles, Profile 1 (mostly permissive) and Profile 6 (permissive, but not frequent app installers) are closely related to the "privacy unconcerned" segment; whereas Profile 2 (mostly protective) and Profile 7 (protective when configuring settings about location) can be seen as relatively more protective, similar to the users in the "privacy fundamentalist" segment. For other profiles we generated, we see quite diverse patterns that target the denial permission decisions on specific types of permissions and categories of apps.
- Lin et al.[59] conducted an online MTurk survey asking participants about their comfort

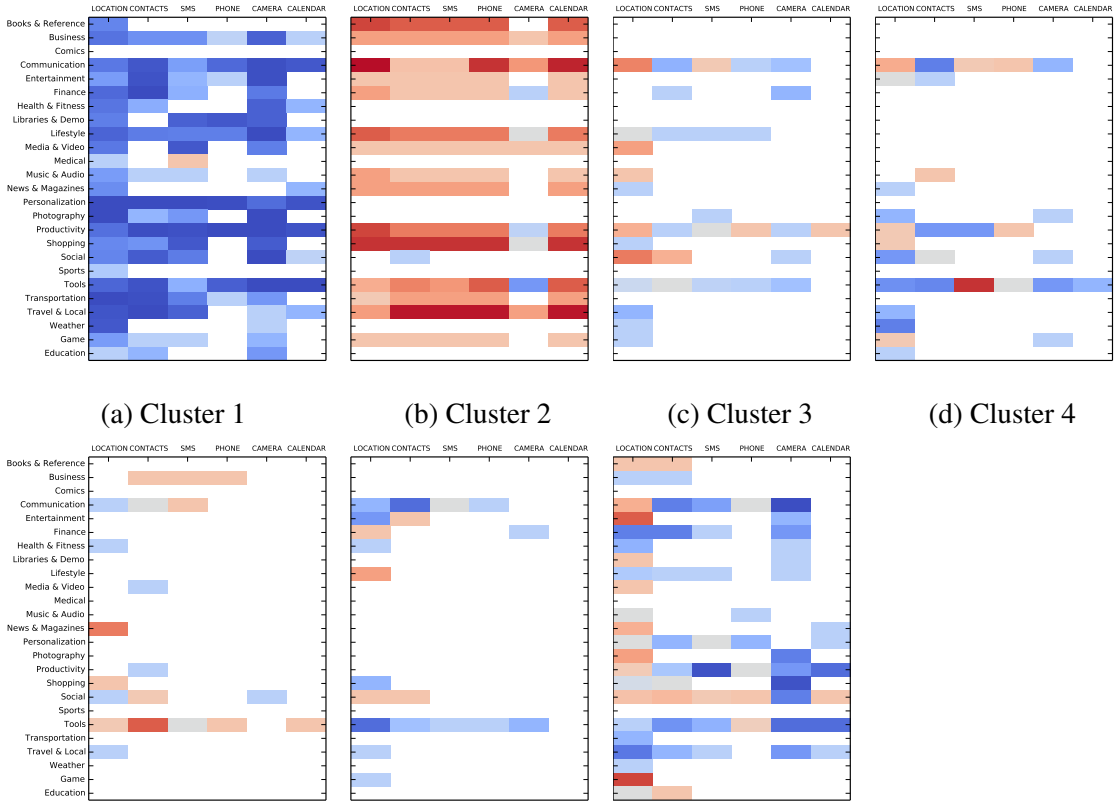

(a) Cluster 1 (b) Cluster 2 (c) Cluster 3 (d) Cluster 4

(e) Cluster 5 (f) Cluster 6 (g) Cluster 7

Figure 4.4: Depiction of Users' Permission Settings Organized by App Categories and Permissions. The colors of the cells indicate how likely users in the cluster are to grant (blue) or deny (red) different permissions to app in different categories. And empty entries (white) indicate that we did not collect any setting data from users in this cluster about this permission and app category.

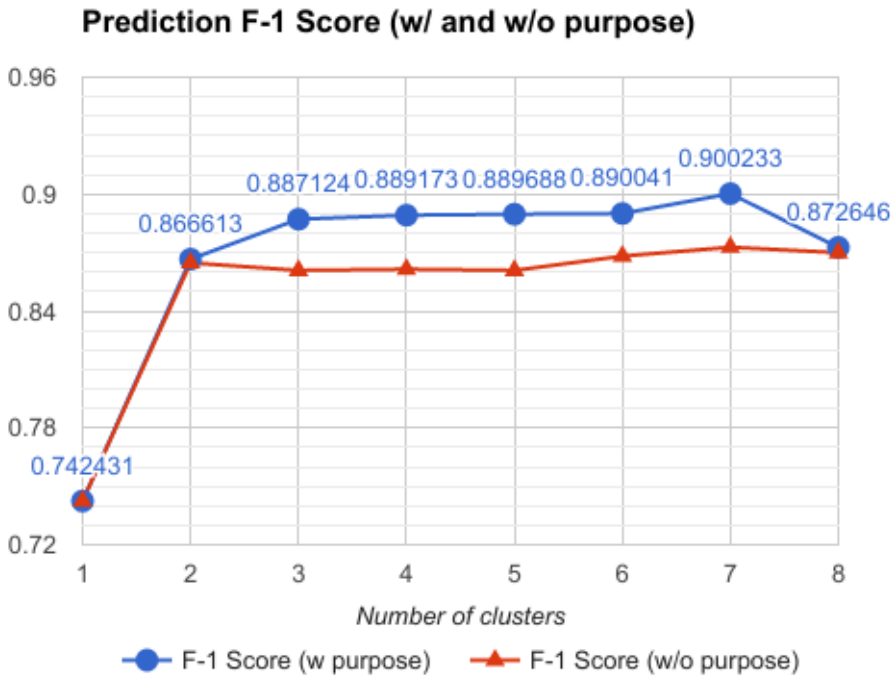

Figure 4.5: Cross-validated F-1 Scores When Predicting Users' Permission Settings Using the Cluster Information Together With App Categories, Permissions, and Purposes. For each choice of K, the number of clusters to generate, we applied a grid search of parameters of both clustering and prediction of users' settings (using SVM classifier). Each number showed the best result of the corresponding number K.

towards allowing a specific app permission request, given the app description and purposes for using the permission data. They segmented users based on aggregations of their responses and categorized them into four profiles: "unconcerned," which is comparable to Profile 1 (mostly permissive); "conservatives," which is comparable to Profile 2 (mostly protective); "fence-sitters," which is comparable to Profile 6 (permissive, but not frequent app installers); and "advanced users," which is comparable to the rest of the profiles in which users generally had mixed preferences or specific concerns on some permissions and categories of apps. One thing to notice is that in the work of Lin et al., the purpose information for the permission requests is completely available, whereas in our study the purpose information is only available for a small fraction (8.6% for location, 35.1% for contacts) of the permission requests. With more purpose information analyzed or required to be shown, we could potentially discover more related patterns between the two sets of profiles.

The evaluation of performance or quality of clustering is unavoidably subjective. People may optimize or tune the clustering parameters for different purposes such as stability, simplicity, and difficulty to interpret. To prepare for generating recommendations on privacy settings for users in a field study in Chapter 5, we tuned the clustering results in this chapter to simply optimize for cross-validated prediction accuracy of profiles generated for each cluster of users. When applying this methodology to other scenarios, one can make different tradeoffs, such as limiting the number of profiles for simplicity or drawing more separable boundaries between clusters, and so on.

To date, neither iOS nor Android has incorporated purpose information into their permission mechanisms. The latest versions of Android (Version 6 or higher) and iOS (Version 7 or higher) provide guidelines and suggestions for app developers to explain their permission requests. However, no terminology or unified categorization has yet been introduced. It remains a challenging and developing research task to infer the purpose information automatically by analyzing the execution logic of the apps. The purpose information we fetched from static analysis of the Android app apk files[59] did not have complete coverage of the apps used by participants in the study (8.6% for location, 35.1% for contact, 42.5% for camera). We would expect better user preference modeling if major mobile operating systems could evolve from purely permission-based mechanisms into a purpose-oriented app permission management.

Because of the limitation of population size, our analysis results were in an anecdotal form. With more user data collected from future studies and the potential adoption of our enhanced permission manager to smartphone operating systems, we could collect more data so that we could derive privacy profiles and profile-based recommendations with higher statistical power. Our down-sampling analysis did show that with a relatively small number of users, we could build a profile-based recommender with predictive power close to models from personalized learning. We believe that with the broader adoption of our proposed technique, the privacy profiles could be interpreted in richer statistical comparisons and frequencies so that the users

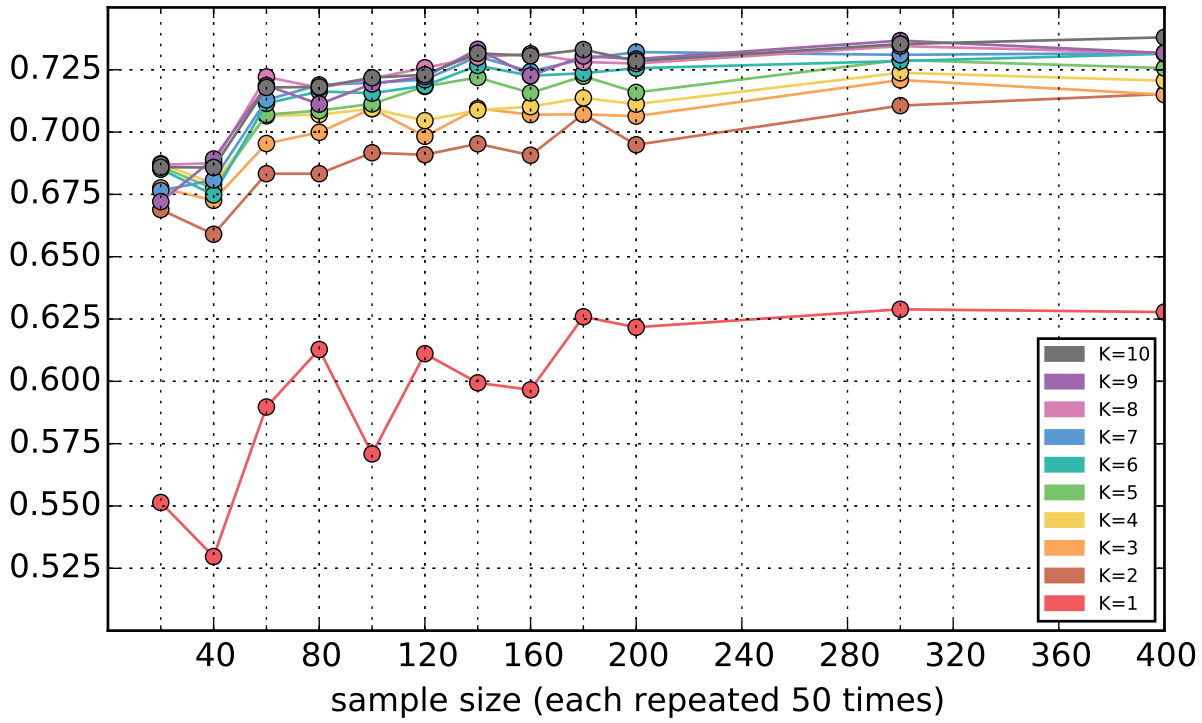

Figure 4.6: Down-sampling Simulation on Lin et al.'s Dataset[59] (F-1 score). With five profiles or more training on data from just 80 users provides a reasonable F-1 score (>0.7). When training on 400 users, the accuracy improves, but only marginally.

could better understand the back-end models of the recommenders for permission settings.

4.4.2 Predicting Users' App Permission Settings Using Profiles

After capturing users' diverse preferences into a small number of profiles, we reduce the task of decision support on users' app permission settings as a profile-oriented classification problem. Given a new user, we first ask the user a small number of questions regarding their app privacy preferences. Using the responses we then estimate the cluster that comes closest to this user in terms of app privacy preferences. Then for each permission request, we can predict the user's decision according to the profile we learned from the cluster of users.

Specifically, in this study, the profile-based recommended settings are generated by a linear-kernel SVM Classifier (LibLinear[34]) on the decision of each permission request. The features of the classifier consist of the user's assigned profile, the category of the corresponding app, the permission requested, and the likely purpose(s) of the permission request. The classifier is pre-trained using the permission settings data we collected when building privacy profiles, with the profile assignment information of the users in the dataset.

Given the relatively small number of 84 participants in our dataset, a potential concern is whether our profiles are expressive enough to cover the privacy preferences of a larger user population and whether we can provide useful recommendations. To explore the utility of our profiles, we conducted an analysis of a larger-scale dataset and down-sampled the subset of data for training the predictive model. We compared the predictive powers of smaller and larger samples.

We applied our approach to building profiles to Lin et al.'s considerably larger dataset[59]. This dataset has 21,657 records in total, consisting of 725 Amazon MTurkers' self-reported preferences of 540 apps accessing permissions for specific purposes, whereas our dataset consists of 3,559 permission settings by 84 participants for 729 apps. To compare the effects of different dataset sizes, we down-sample their dataset by removing randomly selected users to create smaller datasets, ranging from 20 to 400 users in size, which is more than half of the entire dataset. Figure 4.6 shows F-1 scores for 1 to 10 profiles.

The results show that with as little as 80-100 users, which corresponds to our sample size (n=84), the F-1 score can already reach 0.725, only slightly different from the larger sample sizes, which get the best F-1 scores around 0.73. Obviously, with training data from more users, our recommendation accuracy is likely to increase, but this experiment suggests that learning profiles from 84 participants already results in profiles sufficiently stable to be used in practical applications. We expect our methodology could be further extended to have more factors and data in the analysis, as more resources and user study datasets becoming available in the future.

4.5 Discussion

4.5.1 Limitations of Privacy Profiles

First, we discuss the limitations of the data collection study conducted in this chapter.

- **Sample Population.** In this study, the target population available for recruitment was limited because we required root access on participants' devices to apply the permission settings directly for users on their phones. As a result, the sample was skewed toward young, male, tech-savvy, and privacy-conscious participants. Accordingly, one may expect the privacy settings and permission profiles obtained for this population to be more conservative (namely, more restrictive) than those of the general population. Despite this possible limitation, our study led to the identification of diverse profiles, which seemed generally comparable to those identified by Lin in her earlier research[59].

  It is important to understand that the objective of this work was not to identify the "ultimate" privacy profiles for the general population. Rather, our primary objective was to (1) evaluate a practical approach for collecting permission data and learning profiles, and

(2) provide a method for using the resulting profiles in the context of personalized privacy assistants. The work presented herein is particularly important because it relies on the collection of permission data and the validation of personalized privacy assistants in field studies, in which participants used their regular phones in their daily activities. We used nudges to increase participants' awareness and motivate them to engage with the management of their settings. A similar study could be conducted with other target populations, including the general population, given the ability to reliably collect and manage privacy settings on non-rooted phones. Developers who have access to the necessary functionality (whether on smartphones or in other contexts, such as a web browser or a permission manager for a social network) could leverage our approach to learn profiles and provide their users with personalized privacy recommendations. Mobile platform providers, such as Google, Apple, and Samsung, could implement our approach (or provide APIs for researchers and developers) and support functionality similar to the one evaluated in this study.

- **Scale of the Study.** In contrast to our work in Chapter 3, we learned privacy profiles from a relatively small dataset, a process which could be viewed as a limitation. While the numbers are small, we applied privacy nudging to ensure that all the participants were motivated to engage with their app permission settings.

  We collected rich, real-world permission data and aggregated the permission settings by three dimensions, namely app category, permissions, and purpose information. The predictive power of our clusters is further supported by results presented in Chapter 5, where we report on a study in which the profiles were used to recommend permission settings to users: the vast majority of our recommendations (78.7%) were accepted by participants.

- **Study Time Length.** A potential limitation is the relatively short length of our study. It is possible that participants may not have fully converged on stable privacy settings. We believe that the likelihood that this was the case is relatively low because of our use of daily privacy nudges. These nudges were effective at getting participants to review and adjust their permission settings. This approach enabled us to elicit permission settings for a large number of apps (729) and permissions (3,559) in a relatively short time from 84 participants. As Almuhimedi et al.[14] found in an online study, the effects of nudges were typically captured by changes made by users to their permission settings within two or three days from the start of the daily nudges. In future work, we plan to explore longitudinal interactions with personalized privacy assistants over longer periods of time and further study continuous privacy decision-making processes.

4.5.2 Generating Privacy Profiles

In this study, we generated profiles based on users' aggregated preferences for allowing or denying permission, grouped by app category and permissions. We experimented with having purpose

information for permission requests in the study app to help participants better make decisions. However, participants cannot configure settings at the purpose level. Some apps could be requesting permission for multiple reasons or uses (e.g., to support both their core functionality as well as advertising). Participants would have to make the decision to allow or deny a permission request by taking all purposes into consideration. Multiple participants reported that they would have liked to deny certain permissions (e.g., location) for specific purposes (e.g., tracking and profiling), but that they could not do so because it would have broken essential features of the application. This issue suggests that current permission models would benefit from allowing users to grant and deny permissions for specific purposes, rather than forcing users to deny or accept the combination of all purposes. While iOS and Android 6.0 support developer-specified purposes in permission requests[85, 88], once access is granted, apps can use the corresponding resource for any purpose. The current permission model also fails for system services, such as Google Play Services, that provide resource access to multiple apps (e.g., location). Because it is unclear how many apps depend on sensitive resources provided by a service such as Google Play Services, it is effectively impossible for users to make meaningful decisions about granting or denying Google Play access to permissions such as location.

A substantial challenge in mobile computing and other domains will be to shift permission models from resource-centric fine-grained access control (e.g., multiple permissions to read, write SMS) to purpose-centric control that better aligns with users' privacy decision making. While these finer-grained models could increase user burden, our research suggests that they may lend themselves to the learning of more powerful predictive models, an approach which in turn could actually help reduce user burden by providing a larger number of more accurate recommendations.

Apps not functioning or crashing were sometimes reported by some participants as a reason for re-allowing permissions. However, the introduction of a selective permission model in Android 6.0 suggests that in the future, most apps will likely continue to work properly even when requested permissions are denied, as is already the case in iOS, because app developers will adapt and add exception handling for denied permissions. We believe that as privacy protection becomes more advocated, app developers and service providers will eventually adopt scenarios where users agree on granting only parts of the information or functionality requests.

Modeling of users' app permission settings could be implemented in alternative ways, such as applying heuristic rules, building factorization models for collaborative filtering, and so on. We adopted the idea of building privacy profiles by clustering similar-minded users, and generated an aggregated profile of preferences for each cluster of users. This decision was partly due to having only limited lab resources to recruit participants to use our study app in their daily lives. Our analysis in Chapter 3 suggested that we could likely further improve the predictive power of our recommendations by evolving new machine learning tools and gaining more accessible user privacy preference data. With a larger dataset obtainable, we can improve our clustering of users with more data support in each dimension, which will result in a better model with stronger

prediction power fir permission settings of users in each cluster.

4.6 Summary

In this chapter, we conducted a field study to collect Android users' real-world app permission settings. To elicit users' privacy preferences in a more efficient and accurate manner, we collected users' permission settings that reflect the purpose of the permission requests by apps. To make up for the relatively small set of objects, we relied on daily privacy nudges to encourage participants to engage with the permission manager and review/revise permission settings within the study period.

Results of the analysis demonstrated that even with a small number of users, it is possible to build privacy profiles with strong predictive power. We conducted a comprehensive analysis of the privacy profiles generated from the clustering analysis. We compared the profiles with the segmentation and profiles learned from earlier work[59, 89].

We were motivated by the results of the study to explore ways to apply the profiles to help provide personalized decision support for users on mobile app permissions. Specifically, we were interested in how to capture users' privacy preferences by interacting with users in the permission manager, and how users react to our interactive permission manager tool and recommendations.

Chapter 5

A Profile-Based Privacy Assistant to Help Users Configure Mobile App Permission Settings

5.1 Introduction

In Chapters 3 and 4, we analyzed Android users' app permission settings and showed the potential of using a small number of clusters to predict users' permission decisions.

While encouraging, these results still leave a number of questions unanswered. Ultimately, we would like to see whether it is possible to build a privacy assistant that takes advantage of these clusters to effectively help users configure their mobile app permission settings. In this chapter, we discuss the design and evaluation of one such assistant which we piloted with actual smartphone users.

Specifically, we further extended our enhanced Android permission manager app used in Chapter 4 and built a Personalized Privacy Assistant (PPA) app.

- We chose to provide recommendations for permission settings that users can review and selectively accept or reject rather than using a fully automated approach. This decision was motivated by the fact that our predictions are not always correct and also by a desire to not take away control from the user, as retaining a meaningful sense of agency is a key element of privacy.
- We use profiles we generated from Chapter 4 to provide personalized recommendations. This differs from previous work, which relied on one-size-fits-all permission recommendations based on majority voting [12]. In addition, a primary motivation for the work reported in this chapter was to evaluate acceptance of personalized privacy assistant technology by real users in the wild.

- The Personalized Privacy Assistant app relies on a short interactive phase in which it asks up to five questions to the user to determine which cluster of users most closely matches the user's stated privacy preferences. Following this first step, the PPA app relies on the privacy profile associated with this cluster to identify permission recommendations for the apps present on the user's smartphone. Finally, in the third step, these recommendations are presented to the user in a format that enables the user to review the recommendations and decide which of them to accept. Recommendations accepted by the users are used by the Personalized Privacy Assistant to adjust the user's permission settings.

We conducted a pilot study in which participants installed our PPA app and used it on their regular Android phones for a period of 10 days as they went about their regular everyday activities. The pilot included a treatment condition with 49 users who piloted the PPA app, as described above, including the initial interactive dialog to assign users to a cluster, the personalized recommendation phase to identify recommended permission settings, and a final phase where users are presented with recommendations they can review and selectively accept. In addition, the pilot also included a baseline condition involving 23 users who were provided with a bare-bones version of the assistant app, which only included the enhanced permission manager used in our PPA but did not include a dialog to capture the user's privacy preferences or the generation and presentation of permission recommendations. We discuss the results collected as part of the study. The results of the pilot are encouraging and suggest that PPA functionality can help generate useful recommendations, with our interactive configuration leading to high levels of user satisfaction. In particular, 78.7% of the recommended settings were adopted by users in the treatment condition. In addition, despite a week-long regimen of daily privacy nudges designed to motivate participants to review their app permission settings after accepting the PPA's recommendations, only 5% of participants went back to modify recommendations they had originally accepted. This finding further suggests that participants were not simply nudged into accepting the PPA's recommendations, but that accepted recommendations truly aligned with their desired permission settings. The fact that participants did not accept all the recommendations made by their PPA is also further evidence that people carefully reviewed recommendations and did not blindly accept all recommendations made by the PPA. This finding also seems to confirm that an interactive approach such as one where users are able to review recommendations is important for acceptance of this technology.

5.2 Privacy Assistant App to Provide Recommendations

We developed the enhanced permission manager we used in the field study in Chapter 4. The enhanced permission manager allowed users to have easy access to control the permission settings (Figure 4.1, 4.2). We started from the enhanced permission manager and added additional functionality to the app in order to provide profile-based recommendations. Specifically, we

added:

- **Interactive profile assignment dialog.** For a new user of the PPA app, in order to estimate the profile assignment of the user for personalized recommendations, we showed an interactive questionnaire with up to five dynamically generated questions about their app privacy preferences (see Figure 5.1).
- **List of recommended permission settings to deny.** After estimating the profile assignment of the user, the PPA app scanned through all the installed apps on the phone and provided profile-based recommendations according to users' privacy profile (see Figure 5.2). If the PPA app identified permissions to deny, the app would show a list of recommended permissions to deny, for users to review. Users could see the details of the recommendation list and adjust the settings before applying them.

5.2.1 Interactive Profile Assignment

In Chapter 4, we generated a collection of profiles that capture users' diverse app privacy preferences, and provided profile-oriented recommendations. Given a new user that we did not have any prior knowledge of, there would be potentially two ways of eliciting the user's preferences: either wait until the user committed to a sufficient number of privacy decisions, or directly interact with the user to get input. Considering that app permission settings are relatively more privacy-sensitive than other application domains of recommendations such as online shopping, we chose to interact with the users right at the beginning so that we could provide personalized recommendations right away as starting points for users to configure their app permission settings.

Considering that we used the features app categories, permission, and purpose information to cluster users and generate profiles, we designed the profile-assignment questions to use the same features. Each question elicits a user's preference for granting a permission, or allowing a category of apps accessing a permission, or allowing apps accessing a permission for a specific purpose. Users were asked to provide a "Yes" / "No" response to each question. For each user, the PPA app dynamically generated a decision tree[77] that uses input from a question to determine the next question to ask and eventually assigns the user to one of our privacy profiles. Users were asked five questions at most in order to be assigned to a profile.

- The decision tree was generated based on profile assignments and aggregated preferences from the dataset used to build the privacy profiles. For each user in the dataset, we estimated that the user would answer "Yes" if the majority of users' permission settings were "allow"; "No" if the majority of users' permission settings were "deny"; and N/A if this user did not have any related permission setting or had equal numbers of "allow" and "deny" settings.
- We contextualize the decision tree for each user using familiar apps. The study app gen-

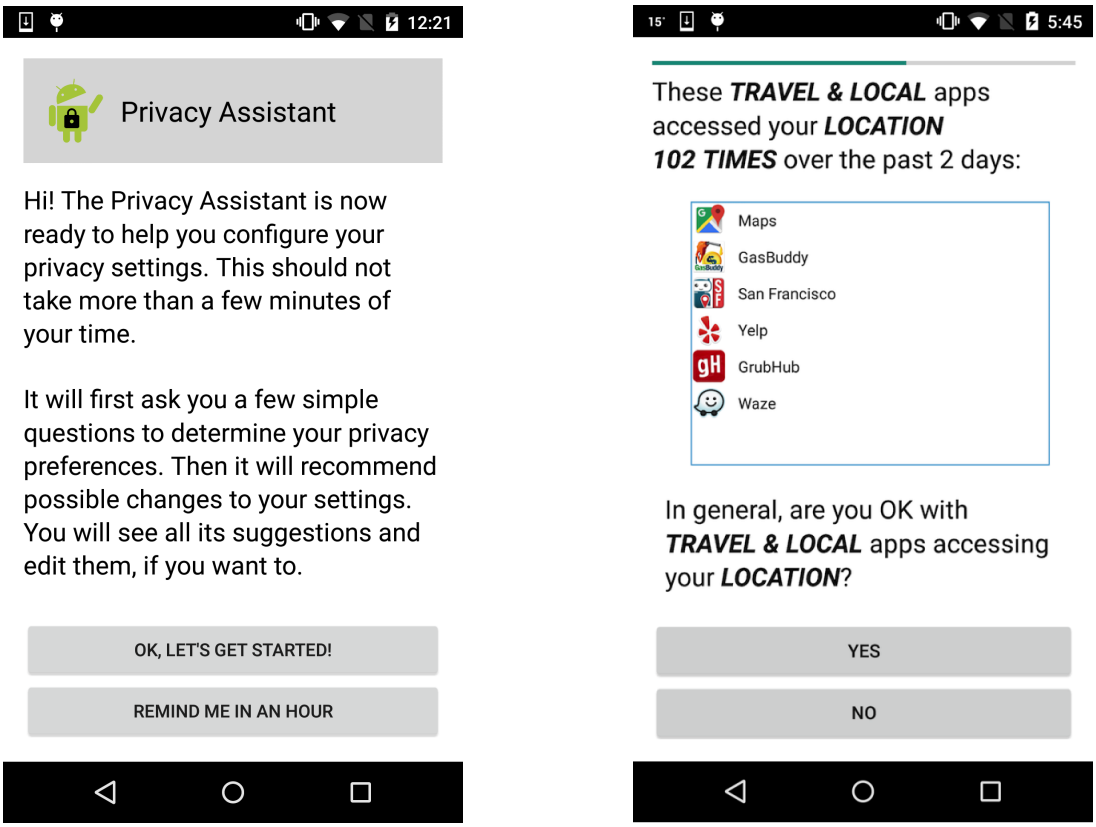

Figure 5.1: Profile Assignment Dialog. Users were asked up to five questions. The sample question on the right asked a user's overall preference for allowing travel & local apps to access location. To further explain the question, we showed access frequency collected on the phone and the list of apps accessing the specific permission.

erates a decision tree dynamically for each user to only include questions related to apps that were installed by this user. For example, if the user had no Game app installed, the PPA would not ask if the user would generally allow Game apps to access location. We excluded questions that were currently not applicable to the user's phone from the process used to generate the tree (for example, if the user had not installed any Game app on the phone, then PPA would not ask if the user would generally allow Location access to Game apps). In this way, the questions could be contextualized using the user's installed apps.

- To further contextualize the questions in the profile assignment dialog, installed apps that fit the particular question were listed in the dialog with their access frequency for the respective permission, inspired by Almuhimedi et al.'s privacy nudges[14]. Figure 5.1 shows an example of an assignment dialog question. In this example, installed apps from the Travel & Local category had accessed the Location permission 102 times over the past two days. A progress bar at the top shows how many questions have been completed.

5.2.2 Showing Profile-Based Recommendations

After receiving the user's responses to the questions, the PPA assigned a privacy profile to the user, which was used to determine which recommendations to show. For each permission requested by apps on the user's phone, the PPA applied the classifier trained with the profiles we generated in Chapter 4 to generate an allow/deny decision for the user. The PPA would then display a list of recommended restrictive permission changes to the user.

Recommendations were grouped by permission (e.g., Calendar, Location); these groups can be expanded to view individual apps, as shown in Figure 5.2. For each app, clicking the question mark reveals an explanation for this specific recommendation, referencing the user's responses to the profile assignment questions. We composed the explanation message with the category of the app and the purpose(s) of this app accessing this permission. For instance, in Figure 5.2 the explanation for denying Snapchat location access is shown. The user can review and adjust recommendation settings. With toggle buttons users can selectively "allow" specific permissions the PPA suggested to deny. The user can accept all shown recommendations, accept some of them by making selective changes, or reject all recommendations.

Thus, based on the privacy profiles generated from real users' privacy settings, our personalized privacy assistant can assign a new user to one of those profiles based on their responses to the profile-assignment dialog. Once a user has been assigned to a profile, we generate recommendations about which permissions a user may want to restrict, personalized to the user's installed apps, by using a classifier with the input of the user's profile and the apps' characteristics, such as its category and the purpose of permission requests.

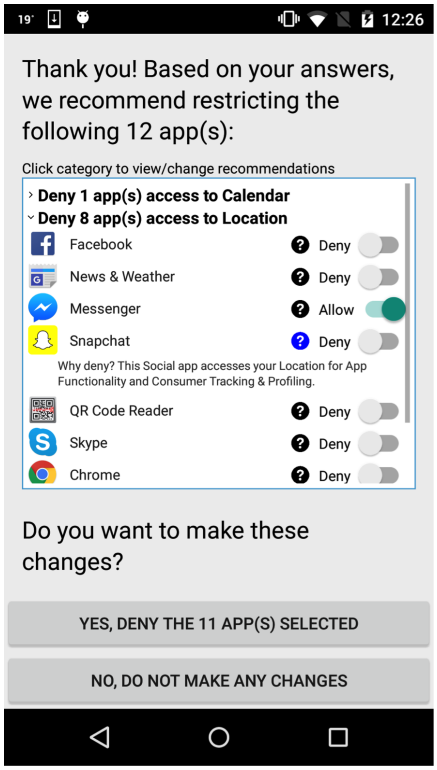

Figure 5.2: Profile-Based Recommendations. After answering up to five questions (see Figure 5.1) users receive personalized recommendations. Users can review and customize the recommended deny settings. In this example, the user re-allowed the Facebook Messenger app to access the user's location after reviewing our initial recommendations. In this particular scenario, the user received recommendations to deny permissions used by 12 apps on the phone. The user reviewed the recommendations and chose to change permission requests by one app back to allow. As the user clicked Yes, the accepted recommendations, which included the deny decisions for the 11 apps listed, were applied to the user's phone.

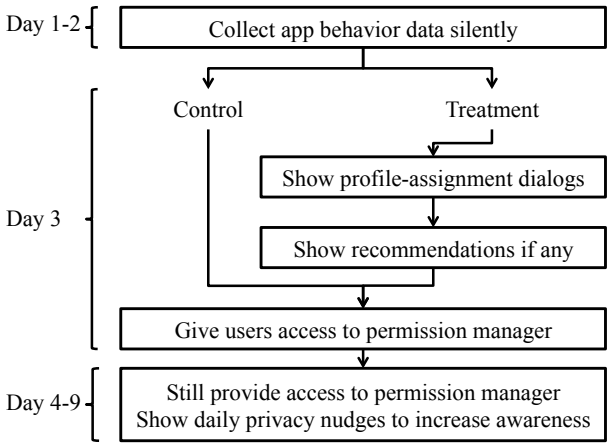

Figure 5.3: Overview of the Study Protocol for the Two Conditions.

5.3 Field Study: Evaluating the App Permission Recommendations

Equipped with the PPA app we developed for profile-based personalized recommendations for app permission settings, we were ready to try the app with real Android users, to see how users interact with the app and manage their app permission settings. We therefore conducted a field study to evaluate our PPA app.

We collected empirical data on how participants interacted with our PPA app and how they modified their permission settings. The study was conducted as a between-subjects experiment with two conditions: (a) the treatment condition in which participants interacted with the PPA, including profile assignment and recommendations; and (b) a control condition without the functionality of profile-based recommendations. Participants in both conditions had access to our enhanced permission manager and received privacy nudges.

5.3.1 Study Procedure

In this field study, our recommendations were generated using the dataset and the profiles we generated from Chapter 4. Thus, we followed the same recruitment approach as for the dataset. We extended the screening survey to exclude those participants who had participated in the study in Chapter 4. To potentially reduce training bias, we also excluded participants with prior experience using other Android permission tools or privacy managing apps. After qualifying for the study, the newly recruited participants received a user ID and instructions for installing the PPA app. Our study protocol is summarized in Figure 5.3.

During days 1 and 2 of the study, the PPA silently collected permission access frequency

statistics for installed apps. Participants did not have access to the permission manager during that time. On the third day, the PPA initiated a dialog with participants.

- In the treatment condition, the app showed an introduction screen and then initiated the profile assignment dialog, in which participants were asked up to five questions about their privacy preferences. Users were then assigned to a profile and personalized recommendations were generated according to the profile assignment. If recommendations could be made, the recommendation screen was shown, and if the PPA did not recommend any changes (i.e., the user was assigned to profile 3), the user was presented with a message saying that it was recommended to keep the current permission settings. The user could review the recommended permission changes and make adjustments as needed. After accepting all, some, or none of the recommendations, participants were asked to rate how comfortable they were with the recommendations on a 7-point Likert scale, followed by a question on why they accepted all, some, or none of the recommendations. After the recommendations and follow-up questions, the PPA opened our permission manager to allow participants to further revise their permission settings.
- In the control condition, the app only showed an introduction screen explaining that users could now change their settings, followed by opening our permission manager. This way, the control and treatment conditions were identical in all aspects, except for the omission of the profile assignment dialog and permission recommendations in the control condition.

Starting on day 4, participants in both conditions started receiving one privacy nudge per day for six days, following exactly the same approach as in the first field study. The goal was to get users to reflect on their privacy settings and thus evaluate whether the profiles matched their preferences or if they would choose to make additional restrictive changes or re-allow any permissions that were restricted based on recommendations. During this phase, we used probabilistic experience sampling (ESM) with single-question dialogs in order to better understand why they denied or allowed permissions, or closed the permission manager without making changes (see Figure 5.4). ESM enabled us to elicit responses from a wider range of participants than would typically agree to participate in exit interviews. ESM dialogs were always consistent with a participant's prior action (e.g., denying permissions). They were shown with 0.66 probability after a user action, to avoid overwhelming users with too many additional dialogs.

At the end of the study, participants were asked to complete an exit survey, which focused on their experience with the profile assignment dialog, perception of the received recommendations, and utility of the additional nudges. After completing the survey, participants were issued a $15 gift certificate. The study received IRB approval.

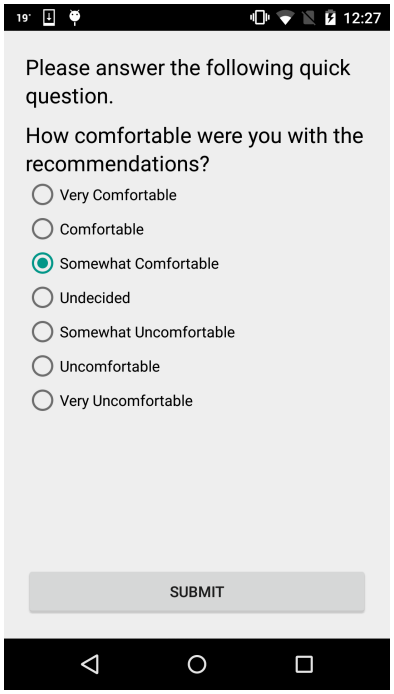

(a) Comfort with the recommendations shown.

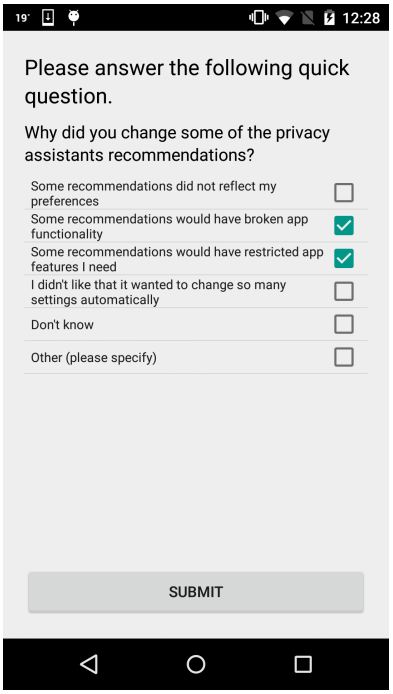

(b) Reason(s) for adjusting the recommendations.

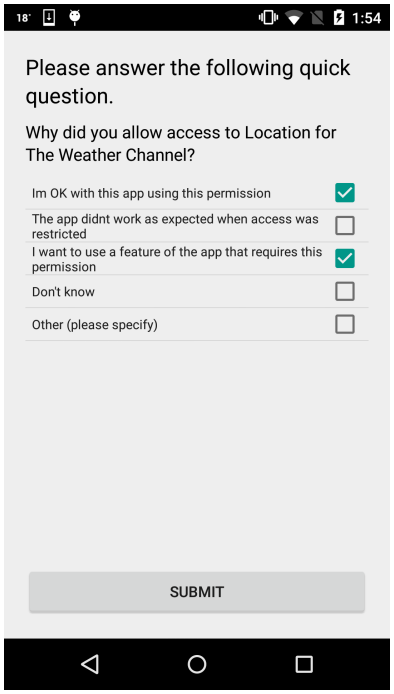

(c) Reason(s) for allowing (or denying) a permission request.

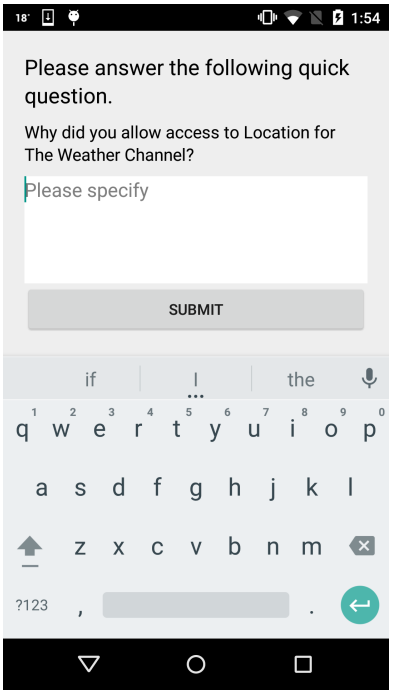

(d) A sample open-ended response. This screen is triggered if a participant clicks "other" in screen (c).

Figure 5.4: Questions Shown to Participants during the Study Using Probabilistic Experience Sampling Method (ESM).

5.3.2 Study Results

We received valid screening survey responses from 138 participants. We excluded four participants who had participated in the first study and three participants who had prior experience with another app privacy manager. Of 131 initial participants, 72 successfully completed the study (49 treatment, 23 control). Participants were randomly assigned to the two conditions in a 2:1 ratio. As the data collection in Chapter 4 suggested, many participants may have permissive privacy attitudes, in which case they may be assigned to the permissive profiles and thus would not receive restrictive recommendations, and hence would not interact with the recommendation screen (shown in Figure 5.2). Thus, we increased the number of treatment participants to account for these considerations.

Our sample population was recruited from the same population as for the data collection study and exhibited similar characteristics. Most participants were male (66 male, 5 female, 1 did not disclose) and originated from North America (56, 52 United States), Europe (7), South America (3), and Asia (2). Among them, 5 had graduate, 17 Bachelor, and 4 Associates degrees; 23 attended some college, 23 had a high school degree or lower. Commonly reported occupations were student (37), computer engineer or IT professional (12), engineer in other fields (6), service (5), and unemployed (3). Participants in this study also exhibited high privacy concerns (IUIPC[66]): control (mean 6.33, median 6, min 4, max 7), awareness (mean 6.67, median 7, min 5, max 7), and collection (mean 6, median 7, min 2.33, max 7).

Effectiveness of Recommendations

As described in Figure 5.3, the participants in the treatment group were shown profile-assignment dialogs followed by potential recommendations. In the treatment group, the number of received recommendations depended on the privacy profile participants were assigned to and their installed apps. Our system only showed recommendations to deny permission requests. For example, if a participant answered mostly "YES" to most profile assignment questions, the system would estimate the user related to a permissive profile and give few recommendations to deny permissions. Likewise, if among the apps installed by the user, if none of them were estimated to be denied by the recommender, the user would not be shown any recommendation. Of the 49 participants in the treatment group, 22 were recommended to keep their current settings. Among them 21 answered "YES" (allow) to most profile assignment questions and were assigned to Profile 1, the most permissive profile. Another participant was assigned to Profile 4 but did not have any of the apps installed that were denied in the assigned privacy profile.

The majority of recommendations were accepted. The 27 participants who received recommendations to deny certain permissions accepted 196 out of 249 individual app recommendations provided (78.7%). Of the 27 participants, 15 accepted all recommendations (4 participants were from profile 1; 3 from profile 2; 6 from profile 3; and 2 from profile 7), 9 accepted some

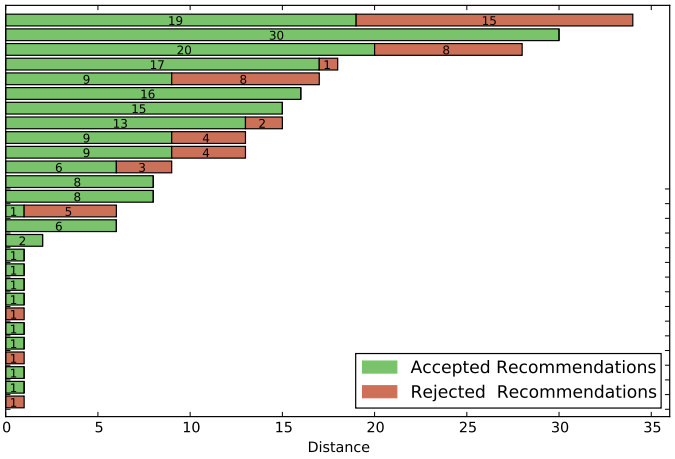

Figure 5.5: Number of Recommendations Accepted or Rejected by Participants Receiving Them. Overall, users accepted 78.7% of all recommendations.

recommendations (2 were from profile 1; 2 from profile 2; 3 from profile 5; and 2 from profile 7), and 3 accepted none (all from profile 3; they were shown only one recommendation).

Figure 5.5 shows the number of accepted and rejected recommendations for each of these participants. To further demonstrate the results, we show the detailed results collected from these 49 participants in the treatment group in Table 5.1 and Table 5.2.

Table 5.1: Detailed Numbers of Participants' Responses and Actions in the Treatment Condition (N=49)

Participant ID	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	T12	T13	T14	T15
Cluster assigned by app dialog before generating recommendations	C1	C4	C3	C3	C5	C1	C3	C1	C5	C4	C1	C1	C4	C1	C1
Cluster estimated from user settings at the end of the field study	C1	C1	C1	C1	C1	C1	C4	C1	C4	C4	C1	C1	C1	C1	C1
Total number of permission settings	124	84	78	79	71	65	86	117	45	53	83	86	67	92	63
Number of deny recommendations	1	13	6	6	34	0	18	1	13	1	0	0	9	0	0
Number of deny recommendations accepted by user	0	9	6	1	19	0	17	1	9	1	0	0	6	0	0
Comfort with the recommendations (7-Likert: 7=very comfortable, 1=very uncomfortable)	7	6	6	6	7	7	6	7	6	6	6	7	5	6	-
Number of deny recommendations accepted by user, but later changed the setting to allow	0	2	1	0	0	0	0	0	0	0	0	0	0	0	0
Number of deny recommendations rejected by user, but later changed the setting to deny	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
Other permissions denied manually by user	0	0	3	0	1	0	4	5	2	0	1	5	12	0	0

Participant ID	T16	T17	T18	T19	T20	T21	T22	T23	T24	T25	T26	T27	T28
Cluster assigned by app dialog before generating recommendations	C7	C1	C1	C1	C7	C4	C1	C3	C3	C1	C1	C1	C1
Cluster estimated from user settings at the end of the field study	C1	C1	C1	C1	C7	C1	C1	C1	C4	C1	C1	C1	C4
Total number of permission settings	125	72	73	113	91	24	39	77	76	96	97	112	69
Number of deny recommendations	17	1	0	1	15	0	0	16	15	0	0	0	0
Number of deny recommendations accepted by user	9	0	0	1	13	0	0	16	15	0	0	0	0
Comfort with the recommendations (7-Likert: 7=very comfortable, 1=very uncomfortable)	5	3	7	7	6	7	5	7	-	6	5	6	7
Number of deny recommendations accepted by user, but later changed the setting to allow	1	0	0	0	0	0	0	1	1	0	0	0	0
Number of deny recommendations rejected by user, but later changed the setting to deny	0	0	0	0	0	0	0	0	0	0	0	0	0
Other permissions denied manually by user	16	8	8	4	4	0	3	4	6	8	1	10	5

Participants kept most of the accepted recommendations.

During the remaining six days of the study after the recommendation dialog (days 4–9), we showed daily privacy nudges to remind users of actual app permission accesses to increase their awareness and engagement. However, only 10 of the previously accepted recommended permission restrictions (5.10% of all accepted recommendations) were re-allowed. This indicates that the privacy choices made based on the recommendations tended to be accurate, and hence

Table 5.2: Detailed Numbers of Participants' Responses and Actions in the Treatment Condition (N=49) (cont.)

Participant ID	T29	T30	T31	T32	T33	T34	T35	T36	T37	T38	T39
Cluster assigned by app dialog before generating recommendations	C1	C1	C1	C1	C7	C1	C1	C1	C1	C1	C1
Cluster estimated from user settings at the end of the field study	C1	C1	C1	C1	C7	C1	C1	C5	C1	C1	C1
Total number of permission settings	87	68	51	25	136	122	131	47	85	54	29
Number of deny recommendations	0	0	1	0	30	1	0	0	0	1	0
Number of deny recommendations accepted by user	0	0	1	0	30	0	0	0	0	1	0
Comfort with the recommendations (7-Likert: 7=very comfortable, 1=very uncomfortable)	7	7	6	7	7	6	7	5	6	4	4
Number of deny recommendations accepted by user, but later changed the setting to allow	0	0	0	0	0	0	0	0	0	0	0
Number of deny recommendations rejected by user, but later changed the setting to deny	0	0	0	0	0	0	0	0	0	0	0
Other permissions denied manually by user	0	1	7	0	0	0	4	0	7	0	3

Participant ID	T40	T41	T42	T43	T44	T45	T46	T47	T48	T49
Cluster assigned by app dialog before generating recommendations	C7	C4	C1	C5	C3	C1	C4	C1	C1	C1
Cluster estimated from user settings at the end of the field study	C3	C1	C1	C6	C4	C1	C1	C1	C1	C1
Total number of permission settings	71	66	75	59	48	58	60	57	138	102
Number of deny recommendations	8	8	0	28	1	1	1	0	2	0
Number of deny recommendations accepted by user	8	8	0	20	1	1	1	0	2	0
Comfort with the recommendations (7-Likert: 7=very comfortable, 1=very uncomfortable)	6	4	7	7	7	7	6	6	7	6
Number of deny recommendations accepted by user, but later changed the setting to allow	1	1	0	1	0	0	0	0	1	0
Number of deny recommendations rejected by user, but later changed the setting to deny	0	0	0	0	0	0	0	0	0	0
Other permissions denied manually by user	0	1	1	0	5	9	4	1	0	11

the recommendations were effective.

Recommendations helped users converge more quickly on settings.

The average numbers of permissions changed by participants per day of the study are shown in Figure 5.6. Among the 383 permission settings changes made by the treatment group, the participants made 316 (82.51%) of them on day 3, which is the day they received profile-based recommendations and the first day they had access to the permission manager. In contrast, the

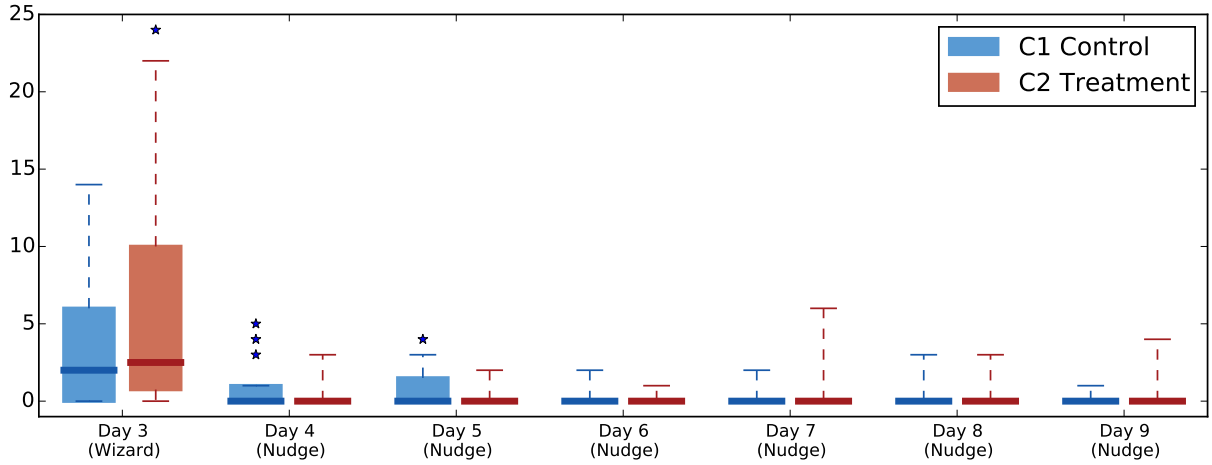

Figure 5.6: Number of Permission Changes in the Control and Treatment Groups on the Different Days of the Study. On day 3, the treatment group received recommendations, and both groups were given access to the permission manager.

control group only made 68.42% (104 of 152) of their permission settings on day 3. The difference between the treatment and control conditions has a significant effect on whether participants made changes on day 3 (logistic regression with user IDs, Odds Ratio=1.72, StdErr.=0.36, z=2.56, p=0.010).

On days 4–9, the treatment group made 67 additional changes to permissions settings (per participant mean 1.39, SD 2.03), and the control group 48 (per participant mean 2.09, SD 2.63). The difference between conditions was not significant. We had 43 related ESM responses from the treatment group and 23 from the control group. Participants gave the following reasons for making restrictive changes: "I don't use the app's features that require this permission" (treatment: 10, control: 6), "I don't want this app to use this permission" (21, 18), "The app doesn't need this permission to function" (16, 11), and "Don't know" (4, 0). This suggests that reasons for restricting permissions were similar across conditions, but the control group had to make more overall changes to arrive at satisfactory settings, whereas the recommendations provided in the treatment group were effective at reducing configuration effort for participants.

In both conditions, a few permissions were restricted and later re-allowed (treatment: 18, mean .62, SD 1.37; control: 11, mean .48, SD .73), with no significant difference between conditions (Mann-Whitney U: U=548.5, z=0.1751, p=0.8572). Participants gave the following reasons for re-allowing: "I want to use a feature of the app that requires this permission" (treatment: 3, control: 1), "I am OK with this app using this permission" (4, 1), "The app didn't work as expected when access was restricted" (2, 1), and "Don't know" (0, 1).

Most participants remain in the same profile they were assigned to.

We collected the participants' app permission settings at the end of the study and compared

them to their responses in the profile-assignment dialogs. For this purpose, we re-ran the profile assignment process with their final permission settings to check their assigned profile, if all their settings were known, and then compared the two assignments for each participant. Of the 49 treatment group participants, 35 (71.43%) remained in the same privacy profile they were assigned to initially. For the other 14 participants (28.57%), their permission settings changes during the study resulted in a different profile being a better fit for them. Two participants switched from profile 1 to profile 2, which generally allows Location access but denies Call Log access. One participant switched from profile 5 to profile 6, which allows more Camera access. One switched from Profile 7 to Profile 1, loosening the restrictions on Social apps. The remaining 10 were re-assigned to Profile 3, which is the most permissive one. A likely explanation is that participants' preferences are more restrictive, but that the lack of ability to control the purposes permissions are granted for forced them to be more permissive than desired, i.e., they lack the capabilities to regulate privacy as desired.

Participants are comfortable with the recommendations provided.

We also collected participants' self-reported comfort with the recommendations and the privacy settings they made during the study. Directly after they accepted recommendations, we asked them to rate their comfort level with the received recommendations on a 7-point Likert scale on the phone screen after they finished interacting with the dialogs. Participants felt very comfortable with the provided recommendations (median 6, mode 7, min 3, max 7).

In the exit survey, we asked participants whether they felt that their permission settings changes during the study had improved their privacy, whether they made all necessary changes, and whether they felt more settings changes were needed. The results are shown in Figure 5.7. We did not find significant differences between the control group and the treatment group (n.s., Mann-Whitney U tests). Participants in both groups felt that their privacy had improved and that they made all the changes necessary for their privacy settings to accurately reflect their privacy preferences. We also did not find significant differences in participants' feelings of a need to make further changes before the settings would reflect their preferences.

Usability of Privacy Assistant App

To evaluate the PPA's usability, we asked Likert-scale and open-response questions to learn what participants found useful or problematic about the PPA, and how it could be improved. We further asked them about the usefulness of the provided recommendations.

Permission manager is useful to monitor apps.

Participants in both conditions stated that they especially liked the ability to monitor apps with our enhanced privacy manager (22 treatment, 12 control). That the PPA was helpful in monitoring apps was also confirmed by treatment group participants when asked about the additional nudges (16). Participants also noted the app's general usability (20 treatment, 11 control).

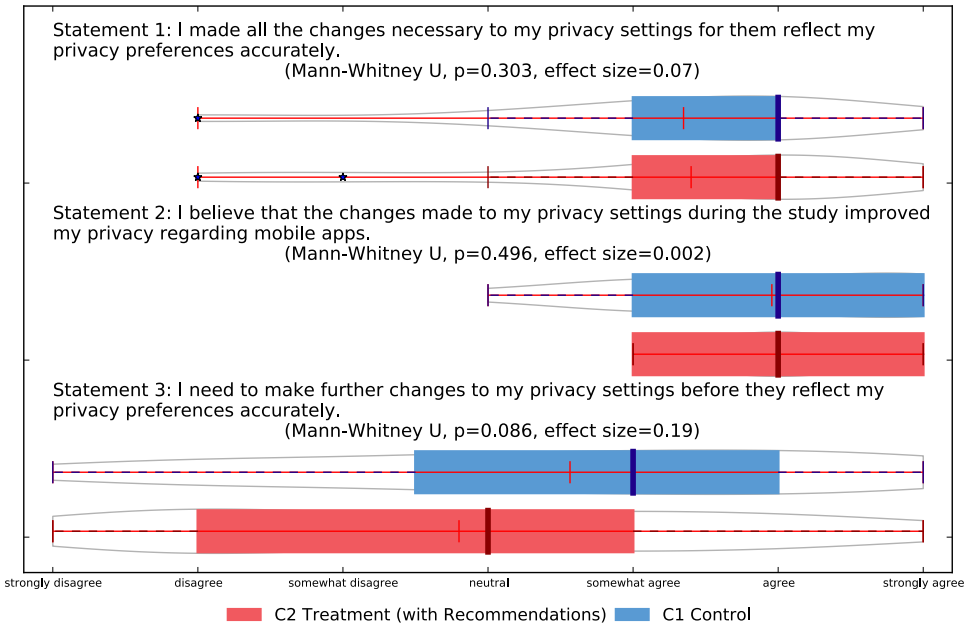

Figure 5.7: Participants' Responses About Their Privacy Settings in the Exit Questionnaire. Participants who received recommendations felt slightly less of a need to make further changes to their settings.

Nudge timing and delivery is important.

When asked about what they liked the least, participants from both conditions identified the timing of the nudges as an issue (18 treatment, 13 control). Asked how we could improve the PPA, participants from both groups suggested turning the nudge into an Android notification (9 treatment, 7 control). Treatment participants also indicated that they would have liked more configuration options (7), mainly to influence the timing of nudges. Note that for study purposes, we purposefully displayed the nudge as a modal dialog to force explicit interaction with the nudge. Finally, it should be stressed that the nudges are not an essential component of the PPA evaluated in this study. They were introduced as part of our empirical protocol to evaluate the stability of settings adopted by participants based on the PPA's recommendations.

Recommendations are helpful.

Of the 49 treatment participants, 27 were shown recommendations, of whom 24 completed the exit survey. Most participants found the recommendations useful (median 5.5, mode 6, min 2, max 7). This was corroborated by free text answers where 13 responses stated that the recommendations provided useful configuration support (11) and decision support (3). P20 stated: "It made what would have taken 10–20 clicks through menus looking to change these settings done in one click," and P10 stated: "It provides you with recommendations using your preferences so you can quickly change the settings without [having] to do much yourself." P4 and P38 found

recommendations useful, but would have preferred to set permissions manually. Four participants found recommendations less useful (3) or useless (1), stating that they prefer to manage settings themselves (1) or that some recommendations would have impaired app functionality (3). Overall, this indicates that recommendations were mostly useful but also points at the issue that users are forced to make tradeoffs when apps crash without permission access. In addition, permissions are currently binary choices: either an app has access to a resource for any purpose or not at all. Restricting permissions for specific purposes is not possible in today's commercial mobile platforms.

Bulk recommendations are useful.

We also asked questions in the exit survey to assess the usability and utility of the different parts of the recommendation screen, such as the timing and amount of information displayed. Participants found that it was useful that all recommendations were listed on one screen (median 6, mode 6, min 3, max 7). This was corroborated by participants disagreeing that it was annoying that they had to click the categories to see details (median 2, mode 2, min 1, max 5). Participants reported their preference for seeing recommendations right after answering each question (median 4, mode 5, min 1, max 6). Participants reported that they somewhat preferred to see the PPA directly after installation (median 5, mode 5, min 3, max 7).

Question dialogs were usable.

Question dialogs were shown to all treatment participants. We asked them to rate on a 7-point Likert scale how easy or difficult the three question types were to answer. All three question types were reported to be easy to answer (permission only: median 7, mode 7, min 3, max 7; permission/purpose: median 6, mode 6, min 3, max 7; permission/category: median 6, mode 7, min 4, max 7). Participants also reported that the app list (median 6, mode 7, min 4, max 7) and access frequency (median 6, mode 6, min 1, max 7) were useful. The app list helped create awareness of how installed apps used permissions (29) and helped to identify apps with undesired permissions (17). Access frequency also helped improve awareness (36) and was mentioned by 6 participants as an important decision factor.

5.4 Discussion

Our pilot study provided rich insights on how users interact with different mobile privacy tools, including our enhanced permission manager, privacy nudge interventions, privacy profile assignment dialog, and profile-based recommendations. Our results show that all these tools play important, yet different, roles in supporting users with privacy configuration and decision-making. As such, these should be taken into consideration when designing personalized privacy assistants and the associated user experience. Results may be even more compelling if we have data from more users and if we have better coverage of purpose information.

Profile assignment is an integral part of our personalized privacy assistant. We use a small number of privacy preference questions to assign users to a profile and provide them with privacy recommendations personalized to their installed apps. Also, we found that participants felt confident answering all three types of questions asked.

Contextualizing the questions with apps that would be affected by the user's response was perceived as useful, and access frequency also helped most users. However, our results indicate that app lists were most helpful in contextualizing profile assignment questions. Current Android and iOS systems do not support stats on the number of times each app accessed a specific permission. However, platform or service providers could get this statistical information by analyzing the frequency of access from other users' phones, so that users could have a better idea on the intensity of permission data access for a new app, even before installing it.

Privacy recommendations introduce a degree of automation to privacy configuration. As Parasuraman et al. pointed out, the degree of automation could potentially impact technology acceptance[72]. Our results indicate that we have achieved a good balance, given that participants reviewed and edited recommendations while reporting high levels of comfort and usability. In future work, we plan to further investigate the impact of different levels of automation on the acceptance of personalized privacy assistants, and the intensity of user involvement for future privacy assistant design.

Our results show that the enhanced privacy manager – including information on both permission access frequency and purpose – helped participants monitor app behavior and manage their privacy settings effectively. Our studies showed that adding the purpose information to a privacy nudge was useful, and nudges were found to be useful in general. Participants liked the utility of frequency and purpose information to help them monitor what apps were doing. A further improvement, motivated by participants' responses, would be to include more information about how privacy and app functionality would be affected by allowing or denying specific permissions.

It would also be interesting to extend our permission manager by adding more privacy options instead of only allowing or denying. For example, one could limit certain app-based features (for example, giving a banking app location access to show nearby branches but not record user location) or purpose-based restrictions (for example, granting Snapchat access to contacts for showing contact names instead of aliases, but restricting access to contacts for user tracking and profiling). However, these additional features would need to be supported by the underlying permission system.

Furthermore, many participants suggested that the design on the timing and modality of the nudge notifications (see a sample screen in Figure 4.3 in Chapter 4) could be refined. This was also raised as an issue by Almuhimedi et al.[14], who applied the nudges to raise awareness of users managing their app privacy settings. However, the use of modal dialog was a conscious choice to force interaction with the nudge messages in our field study. In the public release

version of our PPA, we did not include the pop-up dialog nudge. Non-intrusive messages such as messages and Android push notification are also options to explore.

While our results and insights are primarily related to users' interactions on mobile phones, we expect that personalized privacy assistant approaches can also be applied to support privacy decision-making in other privacy-sensitive domains as well. For instance, websites and online services provide privacy policies that are mostly long and difficult for common users to understand. Another example is the domain of the Internet of Things (IoT), which would involve resolving the privacy preferences of multiple people[82] in a shared/public environment. Also, some IoT ubiquitous smart devices may have small screens or no input access at all. Thus, the IoT permission managers or assistants could be handling multiple devices at the same time. The user interaction design for IoT privacy assistant should be an interesting future research topic.

5.5 Summary

In this chapter, based on the users' privacy profiles that we learned from Chapter 4, we developed a personalized privacy assistant (PPA) app that helps users manage their Android permission settings. Specifically:

- We designed an interactive process that prompts users with a small number of questions so that the assistant app can estimate the profile assignment of users, and thus have a better understanding of the users' privacy preferences. The app then provides profile-based recommendations on permission settings to the users.
- We conducted a pilot study of 72 users (49 treatment, 23 control) who installed our PPA app and used it as a part of a 10-day study on their Android phones. The results showed that the majority of the recommended deny settings were accepted and mostly kept applied to the participants' phones. The recommendations also helped users to converge their settings more quickly. Also, we discussed various design implications from users' behavior and feedback in the pilot study.

We found that users accepted 78.7% of recommendations, and kept 94.9% of recommendations after being shown follow-up daily privacy nudges. Users reported finding the functionality to be beneficial: they generally liked the recommendations and reported a reduction in user burden. We did not have any complaints about the loss of control or autonomy. This indicates that machine learning can help users configure their permission settings by providing recommendations.

Chapter 6

Personalized Privacy Assistant App

In Chapter 3 and Chapter 4, we showed that it is possible to build machine learning models to capture users' app privacy preferences and predict their app permission settings. Then, in Chapter 5 we piloted an assistant app that interacts with the users to provide personalized recommendations on app permission settings to users.

We further enhanced the privacy assistant app and released the app in the Google Play store. In this chapter, we discuss changes made to the design of the Personalized Privacy Assistant (PPA) prior to its release in the Google Play store. This includes improvements to the user interface as well as using additional data collected as part of the study detailed in Chapter 5 to refine the profiles used to recommend permission settings.

6.1 User Interface Design

Prior to releasing our mobile app permission PPA in the Google Play store, several changes were made to the PPA's user interface as well as the model used to make recommendations. These changes are discussed below.

- As reported in Chapter 5, the PPA app that we piloted took two days to observe the behavior of apps installed on the user's phone, collecting permission access frequency statistics. This data was then used in the initial dialog with the user (Figure 5.1) to elicit information about the user's privacy preferences. In contrast, we wanted the version of the PPA deployed in the Google Play Store to be usable from the moment it was downloaded by the user. This implied that we would no longer be able to collect permission access frequency statistics for the apps residing on the user's phone. Instead we had to revise the design of the screens used to support the initial dialog with the user. The end result is a cleaner screen design but also one that is less specific, as we no longer have access to statistics on permission access frequency. As part of this new design, we also opted to give users three

options. Instead of the "yes" or "no" choices offered in the PPA piloted as part of the study detailed in the previous chapter, we introduced a "not sure" option, recognizing that perhaps with less information users would also be less confident about their own preferences (see Figure 6.1 (a)).

This change in turn required adjusting the way we trained the decision tree used to assign people to clusters. When generating the contextualized decision tree in Chapter 5, we assumed that a user would have no answer to a question only if the user had no related permission setting or had equal numbers of "allow" and "deny." In the new version, we label a user's preference in a different way: we label the user's preference as "mostly OK" if more than $T\%$ of the settings were "allow"; "mostly no(t)" OK if more than $T\%$ of the settings were "deny"; and "not sure" for all other cases.

In the first iteration of this PPA app in the Google Play store, we explored different thresholds T indicating a strong opinion about a question. The higher the threshold, the less sensitive the profile-assignment questionnaire is when categorizing users into clusters. We decided on setting a threshold at the point where two-thirds of the settings related to this question were allow or deny. We chose this threshold to give us a good balance of prediction power and data sparsity using the users' data collected from the previous field studies.

- In the field study in Chapter 5, after answering the questions the participants received the recommendations on permissions to deny in batches. After reviewing the list of recommendations, users could apply the adjusted settings in one click ("deny the items selected" or "do not make any changes" (see Figure 5.2)). These bulk recommendations were praised by the participants in the exit survey. However, we did not follow up with providing recommendations for newly installed apps within the study time period. In the new scenario, users would be expected to keep installing and using new apps on their devices.

  In the new PPA app, we apply three changes to the way recommendations are shown: (1) The app no longer shows a batch recommendation screen. Instead, it shows recommendations for items to deny in the permission manager. If the PPA app recommends denying a permission request which is currently allowed, there will be a highlight mark next to the setting button. The mark will disappear if the user toggles the button to deny the permission access manually. (2) If the PPA app detects installation of a new app, it will generate recommendations for this app and show them to the user if it has recommendations for permissions to deny. (3) Users can re-do the profile-assignment questionnaire at any time to refresh the profile assigned to them.

At this time, the PPA app in the Google Play store is only available to users with rooted Android phones, as the standard version of Android does not allow app developers to directly manipulate app permission settings. Instead of collecting usage data in field studies in a mandatory manner, we chose to collect anonymous usage statistics only if the user voluntarily opts-in to share data with us. When a new user opens the PPA app for the first time or when a user clicks

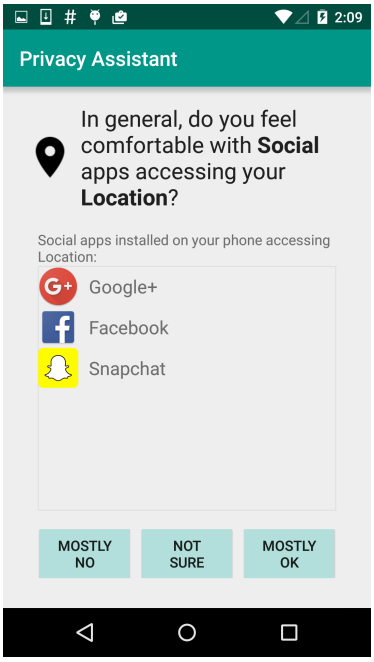

(a) Questions asked by the PPA app

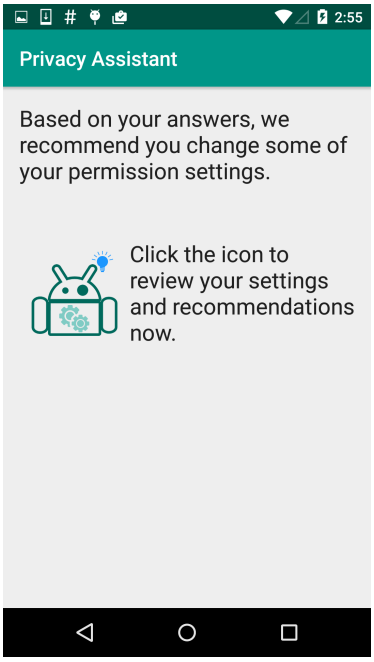

(b) Message shown to the user if PPA finds permission settings it recommends to deny.

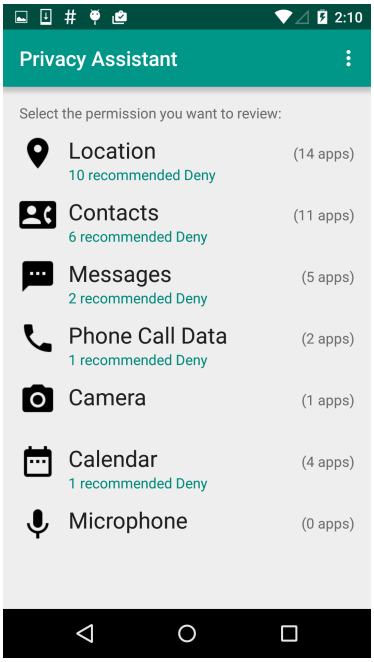

(c) After answering the questions, the user will be navigated to the app home screen where they can review and configure app permission settings.

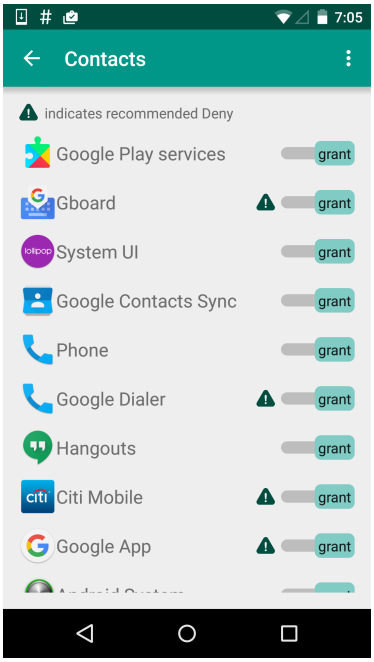

(d) The screen to configure settings for a specific permission (Contacts in this figure) in detail. For an app requesting a permission, if PPA recommends denying but the current setting grants the permission, a highlight mark is shown next to the toggle button.

Figure 6.1: The Personalized Privacy Assistant App on Google Play Store

consent from the menu, the PPA app shows a consent notice to the user and allows the user to choose whether to opt-in to data sharing. Users can opt-out at any time during the use of the app.

We also modified the user interface design of the app to fit the newer appearance of Android operating systems (see Figure 6.1). We also adopted elements from Google's material design guidelines. [1]

6.2 Generating Privacy Profiles and Recommendations

We decided to use the data collected as part of the study reported in Chapter 5 to enhance the profiles generated in Chapter 4. As a result, our dataset of users' settings is from a larger collection of people – a total of 156 participants. Among these participants, we have 84 participants who contributed to the data collection in the field study of Chapter 4, 49 participants who piloted the profile-assignment dialog and recommendations, and 23 people in the control group, namely the group of participants who did not receive recommendations in Chapter 5. We recognize that the data collected to build profiles in Chapter 4 and the data collected from the participants in the two conditions in study described in Chapter 5 are not entirely identical. Nevertheless, the permissions collected for users in all three of these datasets were permission settings configured by users after they were subjected to a similar regimen of daily nudges for at least a week. As we have seen in previous studies, such a regimen of daily nudges has been shown to be very effective at motivating users to carefully revisit and adjust their permission settings[14]. Accordingly, one can reasonably assume that, while some of the earlier steps that subjects in each of the three groups went through were different, the effects of these differences were by and large eliminated by the daily nudges to which participants in each of these groups were subjected. Given the scarcity of available data, we decided that it was reasonable to combine permission settings collected in each of these datasets for the purpose of building privacy profiles for the version of the PPA to be released in the Google Play store. As the discussion below will show, this data did indeed yield a small number of homogeneous clusters with each cluster having a larger number of participants than the clusters used in the previous chapter. This in turn seems to validate our decision to combine data from all three of these datasets for the purpose of generating privacy profiles.

Finally, because Android does not support purpose-specific permission settings and because static analysis does not allow one to systematically infer the purpose(s) for which each app requests a permission, we also chose not to include purpose information in the Google Play version of the PPA app. Instead, we opted to build coarser profiles that do not differentiate between the purpose(s) for which a permission is requested by an app.

Thanks to our a larger dataset, our privacy profiles are able to generate two sets of recom-

[1] https://developer.android.com/guide/topics/ui/look-and-feel/

mendations. When an app is among the 100 most popular apps in our dataset, our new PPA is able to generate app-specific recommendations. For other apps, the PPA's recommendations are based on the app's category, as was the case in the PPA piloted in the previous chapter.

We were also able to compare users' willingness to grant different permissions to the 100 most popular apps, versus granting the same permissions to less popular apps. This comparison shows that users are more likely to allow permission requests when these requests come from the 100 most popular apps ($\chi^2 = 5.6606, df = 1, p = .01735, odds - ratio = 1.1918$). This phenomenon is attributed to a brand effect.

We also observed that users' permission settings were usually not evenly distributed among different app categories and permissions. The top three categories (Communication 21.71%, Tools 18.53%, Productivity 12.25%) covered 52.49% of all permission settings in the dataset. Similarly, the top three permissions (Location 44.83%, Contacts 24.14%, SMS 12.05%) covered 81.02% of settings. Thus, the sparsity of users' decisions could affect the prediction accuracy of the profile assignment. We applied a weight scale to amplify preferences for entries for which we had a larger number of decisions. Specifically, when generating feature vectors for each user prior to clustering them, instead of directly counting the frequency of allows and denies, we used the following formula: for each user u in the dataset, we model the user's preference over permission p for a given category of apps c, where user u has opted to "Allow" that permission for a apps in the given category and "Deny" it for d apps in that category as:

$$P(u, c, p) = \frac{2}{1 + exp(-0.5(a + d))} * \frac{a - d}{a + d}$$

With re-weighting of data using the Sigmoid formula above, we can expect that if two users have the exact same preference, the user with more data points will have bigger value of $P(u, c, p)$. For example, if the user has 20 allows and 5 denies for a given category c and permission p, the preference score $P(u, c, p)$ is 1.2; in contrast, if the user only has 4 allows and 1 deny, the score goes down to 1.109.

We also experimented with other classification algorithms, taking advantage of our larger corpus of permission settings. Instead of using a linear SVM classifier as we had done in Chapter 3 and Chapter 4, we found that a Random Forest classifier resulted in even better scalability and accuracy. Figure 6.2 shows the overall F-1 score of classifier performance for each combination of algorithm and data input. From the figure, we can observe that the Random Forest classifier outperforms the SVM Linear models if we provide separate predictions for the top 100 apps.

6.3 A Close Look at Privacy Profiles

As already discussed in the previous section, we made several changes to both the front-end and back-end functionality of the PPA app prior to publishing it in the Google Play store. Another

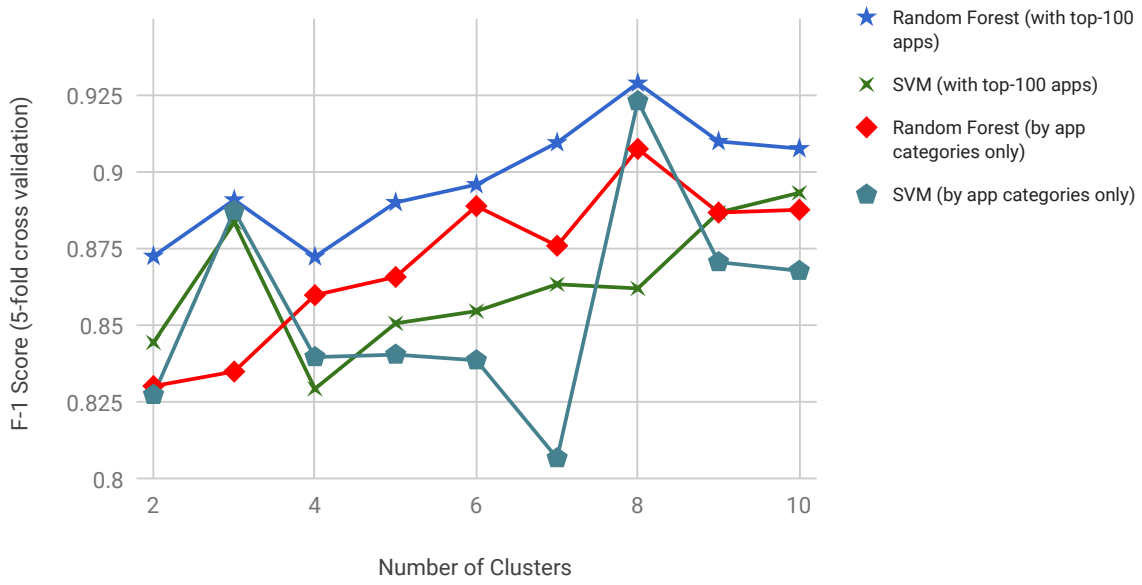

Figure 6.2: Cross-validated F-1 Scores for Permission Recommendations. (PPA profiles, based on dataset from 156 users.) The Random Forest classifier seems to outperform the SVM Linear models; app-specific recommendations for the top 100 most popular apps seem to also have somewhat stronger predictive value than recommendations for entire categories of apps.

change had to do with the criteria used to configure our clustering parameters. As reported in Chapter 4, the clusters built for the PPA piloted in the study reported in the previous chapter were generated using parameters configured to maximize F-1 prediction scores on the training data. In contrast, for the PPA released in the Google Play store, we opted to pick a number of clusters that would offer a good compromise between accuracy and interpretability of the clusters. In contrast to the seven clusters used to organize our set of 84 users in Chapter 4, we ended up with just six clusters for the 156 users available in the larger dataset used to train the PPA released in the Google Play store. The resulting clusters each contain a larger number of individuals and allow for a finer comparison. Specifically, the number of users in each cluster were now 17, 39, 25, 21, 17, and 37. These clusters are the ones used by the PPA we published in the Google Play store to recommend permission settings to users.

Below we take a closer look at these clusters. Each cluster is depicted as a collection of rows corresponding to different app categories or different apps (we show only the 20 most popular apps) and a collection of columns corresponding to six top-level permissions, namely location, contacts, SMS, phone, camera, and calendar. For each cluster, we show two sets of tables:

- Aggregated preference data for users in each cluster when it comes to granting or denying different permissions. In each cluster, we show aggregate preferences for each of the top 20 most popular apps as well as aggregate preferences for apps in each category of apps –

including the 100 most popular ones. Each entry also includes the number of allows and denies recorded for users in the given cluster and a color summarizing these numbers, with shades of blue indicating that users in the cluster lean towards granting the corresponding permission and shades of red indicating that users in the cluster lean towards denying the permission. White is used to denote cells with no data for users in a given cluster. Color depth indicates the level of agreement, with a darker color indicating strong agreement among users in the cluster and a lighter color indicating weaker agreement.

- Permission recommendations indicate how the PPA turns the underlying preference data for users in a given cluster into permission recommendations for users assigned to that cluster. This includes showing separate recommendations for the top 20 most popular apps, as well as recommendations organized by categories for apps other than the 100 most popular apps. The color of each cell indicates the recommendation decision: blue for allowing, red for denying, and white for no recommendation because there is insufficient confidence to make a recommendation. It is important to note that our classifier is trained with five-fold cross-validation to reduce the chances of over-fitting. Thus, some cells will have a red color (deny) even if we observe some allow decisions from users in the dataset.

To further clarify the numbers in each entry, let us consider users in Cluster 1, the app category "Books and References," and the "Location" permission. Overall, users in this cluster have allowed this permission six times and denied it once. If we now consider recommendations for apps that are not part of the 100 most popular apps, we note that the PPA still errs on the safe side and recommends denying this permission based on the fact that, excluding the 100 most popular apps, there are only 4 data points available for users in this cluster: 3 allows and 1 deny. While there are more allows than denies, this is not considered sufficient to recommend allowing this permission to users in the cluster. On the other hand, apps in the next category have a total of 17 data points for users in this cluster (for apps other than the 100 most popular apps), with 16 "allow" and 1 "deny." This is deemed sufficient by the PPA to recommend allowing this permission to all apps (other than the 100 most popular apps) in this category for users in Cluster 1. Finally, consider the Communications category and the Location permission. Overall users in Cluster 1 lean towards granting this permission (93 "allow" and 6 "deny"), but many of the apps in this category are particularly popular and users here show somewhat more nuanced preferences for different popular apps. For instance users in Cluster 1 do not seem as willing to share their location with Facebook Messengers (22 "allow" but 4 "deny") as with Google Chrome (35 "allow" versus 1 "deny"). As a result the PPA recommends that users in this cluster deny this permission to Facebook Messenger (again erring on the safe side) while recommending they grant it to Google Chrome. In addition, when it comes to communications apps that are not among the 100 most popular apps, the PPA recommends actually granting this permission, based on 12 "allow" versus only 1 "deny."

The figures below are ordered by cluster IDs (Figures 6.3, 6.5, 6.7, 6.9, 6.11, and 6.13 for aggregated preferences; Figures 6.4, 6.6, 6.8, 6.10, 6.12, and 6.14 for expected recommendations).

From the results, we can see correlations between our clusters of users and previous segmentation approaches.

- Cluster 1 and Cluster 2 have the most permissive profiles. These users seem similar to the group of people Westin refers to as the "Privacy Unconcerned" Index[55, 93, 94]. Our model split these permissive users into two profiles to further improve the prediction accuracy of our recommendations. Even though they generally have permissive preferences, users in these clusters still get a few "deny" recommendations.
- Cluster 4 has the most protective profile. People in this cluster seem to most closely approximate Westin's "Privacy Fundamentalists" Index[55, 93, 94]. Even though they are likely to have strong privacy concerns, users in this clusters do not receive only "Deny" recommendations. Recommendations also include some "Allow" (e.g., Google Map accessing location data).
- Cluster 3, Cluster 5, and Cluster 6 could be seen as Westin's "Privacy Pragmatists"[55, 93, 94]. They exhibit more nuanced privacy concerns. For example, users in Cluster 5 generally lean towards denying location access more than other permissions. Cluster 5 seems to exhibit similarities to the "Advanced Users" cluster identified Lin et al.[59].

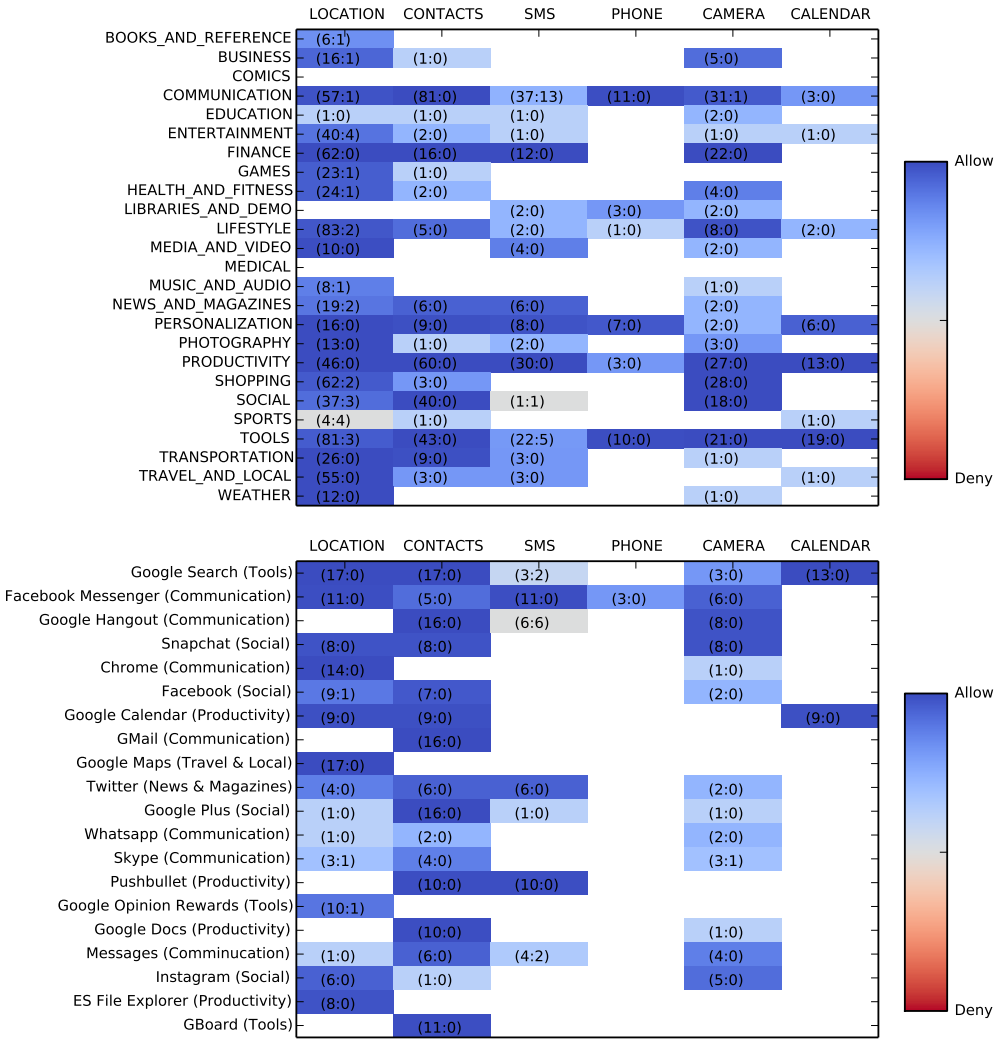

Figure 6.3: Permission Settings of Users in Cluster 1 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

Figure 6.3 shows data for people in Cluster 1, separating permission preferences for the top 20 most popular apps. People in this cluster (17 out of 156) seem to be mostly permissive when it comes to granting permissions to apps. Some of the deny decisions were for sports app accessing location, and various communication apps accessing SMS. One potential explanation for this result would be that users most often use only one app to handle the SMS data instead of managing them across different apps. Note that for some less frequent permission requests, such as the "Chrome" browser accessing the camera, the recommender would choose not to provide any recommendation. Users would be prompted for the Android 6+ version or allow the permission by default for earlier Android versions.

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
BOOKS_AND_REFERENCE	Deny (3:1)	NoRec	NoRec	NoRec	NoRec	NoRec
BUSINESS	Allow(16:1)	Allow(1:0)	NoRec	NoRec	Allow(5:0)	NoRec
COMICS	NoRec	NoRec	NoRec	NoRec	NoRec	NoRec
COMMUNICATION	Allow(14:0)	Allow(16:0)	Allow(8:2)	Allow(5:0)	Allow(3:0)	Allow(0:0)
EDUCATION	Allow(1:0)	Allow(1:0)	Allow(1:0)	NoRec	Allow(2:0)	NoRec
ENTERTAINMENT	Allow(28:1)	Allow(2:0)	Allow(1:0)	NoRec	Allow(1:0)	Allow(1:0)
FINANCE	Allow(39:0)	Allow(2:0)	Allow(2:0)	NoRec	Allow(12:0)	NoRec
GAMES	Allow(23:1)	Allow(1:0)	NoRec	NoRec	NoRec	NoRec
HEALTH_AND_FITNESS	Allow(15:0)	Allow(2:0)	NoRec	NoRec	Allow(1:0)	NoRec
LIBRARIES_AND_DEMO	NoRec	NoRec	Allow(1:0)	Allow(1:0)	Allow(2:0)	NoRec
LIFESTYLE	Allow(76:1)	Allow(5:0)	Allow(2:0)	Deny (1:0)	Allow(6:0)	Allow(2:0)
MEDIA_AND_VIDEO	Allow(10:0)	NoRec	Allow(4:0)	NoRec	Allow(2:0)	NoRec
MEDICAL	NoRec	NoRec	NoRec	NoRec	NoRec	NoRec
MUSIC_AND_AUDIO	Allow(8:1)	NoRec	NoRec	NoRec	Allow(1:0)	NoRec
NEWS_AND_MAGAZINES	Allow(5:0)	Allow(0:0)	Allow(0:0)	NoRec	Allow(0:0)	NoRec
PERSONALIZATION	Allow(14:0)	Allow(7:0)	Allow(6:0)	Allow(6:0)	Allow(2:0)	Allow(6:0)
PHOTOGRAPHY	Allow(3:0)	Allow(0:0)	Allow(0:0)	NoRec	Allow(3:0)	NoRec
PRODUCTIVITY	Allow(11:0)	Allow(10:0)	Allow(7:0)	Allow(0:0)	Allow(17:0)	Allow(3:0)
SHOPPING	Allow(32:1)	Allow(0:0)	NoRec	NoRec	Allow(10:0)	NoRec
SOCIAL	Allow(10:1)	Allow(6:0)	Deny (0:1)	NoRec	Allow(1:0)	NoRec
SPORTS	Deny (4:4)	Allow(1:0)	NoRec	NoRec	NoRec	Allow(1:0)
TOOLS	Allow(22:0)	Allow(9:0)	Allow(10:0)	Allow(5:0)	Allow(10:0)	Allow(6:0)
TRANSPORTATION	Allow(16:0)	Allow(2:0)	Allow(1:0)	NoRec	Allow(0:0)	NoRec
TRAVEL_AND_LOCAL	Allow(32:0)	Allow(0:0)	Allow(0:0)	NoRec	NoRec	Allow(0:0)
WEATHER	Allow(12:0)	NoRec	NoRec	NoRec	Allow(1:0)	NoRec

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
Google Search (Tools)	Allow (17:0)	Allow (17:0)	Deny (3:2)	NoRec	Allow (3:0)	Allow (13:0)
Facebook Messenger (Communication)	Allow (11:0)	Allow (5:0)	Allow (11:0)	Allow (3:0)	Allow (6:0)	NoRec
Google Hangout (Communication)	NoRec	Allow (16:0)	Deny (6:6)	NoRec	Allow (8:0)	NoRec
Snapchat (Social)	Allow (8:0)	Allow (8:0)	NoRec	NoRec	Allow (8:0)	NoRec
Chrome (Communication)	Allow (14:0)	NoRec	NoRec	NoRec	Allow (1:0)	NoRec
Facebook (Social)	Allow (9:1)	Allow (7:0)	NoRec	NoRec	Allow (2:0)	NoRec
Google Calendar (Productivity)	Allow (9:0)	Allow (9:0)	NoRec	NoRec	NoRec	Allow (9:0)
GMail (Communication)	NoRec	Allow (16:0)	NoRec	NoRec	NoRec	NoRec
Google Maps (Travel & Local)	Allow (17:0)	NoRec	NoRec	NoRec	NoRec	NoRec
Twitter (News & Magazines)	Allow (4:0)	Allow (6:0)	Allow (6:0)	NoRec	Allow (2:0)	NoRec
Google Plus (Social)	Allow (1:0)	Allow (16:0)	Allow (1:0)	NoRec	Allow (1:0)	NoRec
Whatsapp (Communication)	Allow (1:0)	Allow (2:0)	NoRec	NoRec	Allow (2:0)	NoRec
Skype (Communication)	Allow (3:1)	Allow (4:0)	NoRec	NoRec	Allow (3:1)	NoRec
Pushbullet (Productivity)	NoRec	Allow (10:0)	Allow (10:0)	NoRec	NoRec	NoRec
Google Opinion Rewards (Tools)	Allow (10:1)	NoRec	NoRec	NoRec	NoRec	NoRec
Google Docs (Productivity)	NoRec	Allow (10:0)	NoRec	NoRec	Allow (1:0)	NoRec
Messages (Communication)	Allow (1:0)	Allow (6:0)	Deny (4:2)	NoRec	Allow (4:0)	NoRec
Instagram (Social)	Allow (6:0)	Allow (1:0)	NoRec	NoRec	Allow (5:0)	NoRec
ES File Explorer (Productivity)	Allow (8:0)	NoRec	NoRec	NoRec	NoRec	NoRec
GBoard (Tools)	NoRec	Allow (11:0)	NoRec	NoRec	NoRec	NoRec

Figure 6.4: Expected Recommendations for Users in Cluster 1 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

Figure 6.4 also shows that even though in this cluster the PPA generally recommends "Allow" for tool apps requesting access to the SMS permission, this does not extend to "Google Search." The PPA actually recommends denying access to SMS by "Google search" (3 "allow" but 2 "deny"). We speculate that the other apps in this category are perceived as having better reasons to request access to this permission than "Google Search."

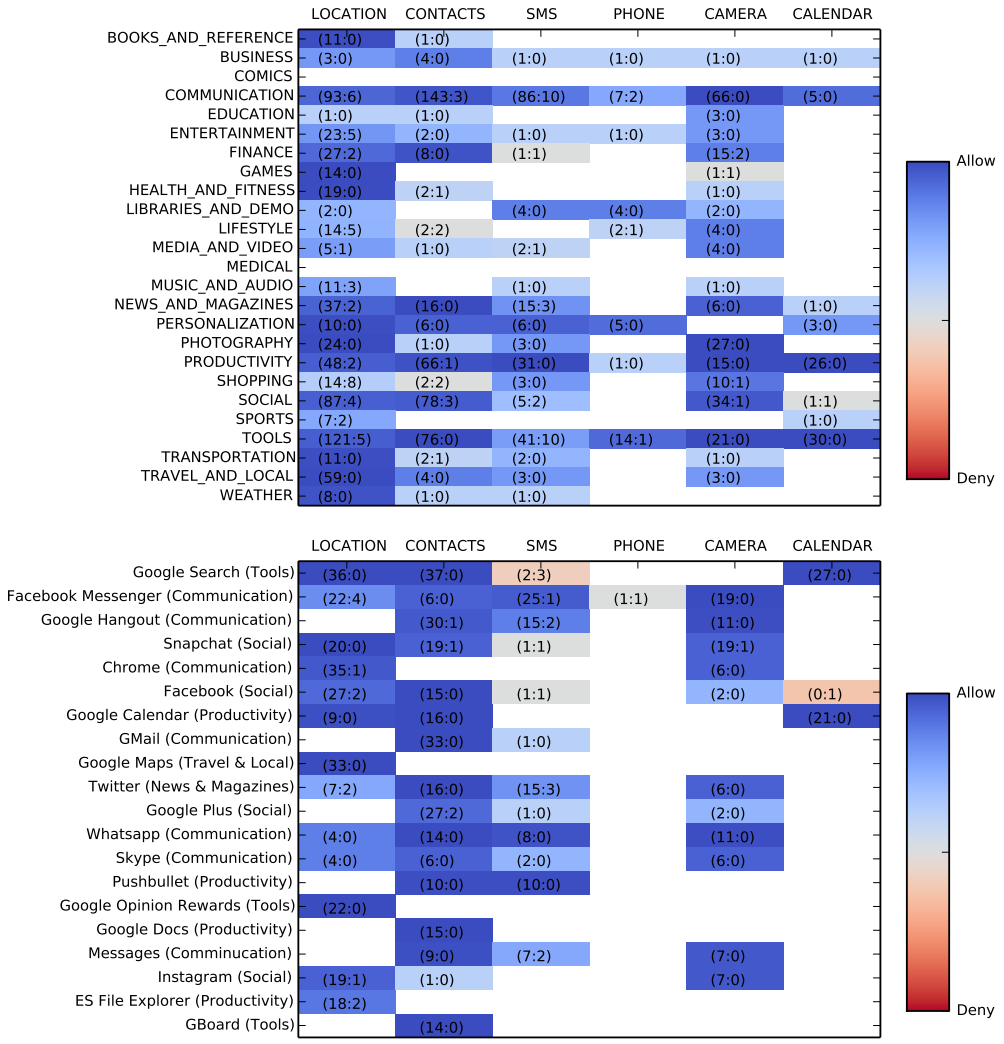

Figure 6.5: Permission Settings of Users in Cluster 2 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

Figure 6.5 shows that users in this cluster (39 out of 156 users) generally share many similarities with users in Cluster 1 and generally have particularly permissive settings. Yet they are more cautious when it comes to granting access to their location and contacts to some apps. We observe a small number of "deny" for shopping apps, entertainment apps, and lifestyle apps requesting access to location. For example, among 26 users using "Facebook Messenger," four denied location access, and two out of three "Walmart" app users denied location access.

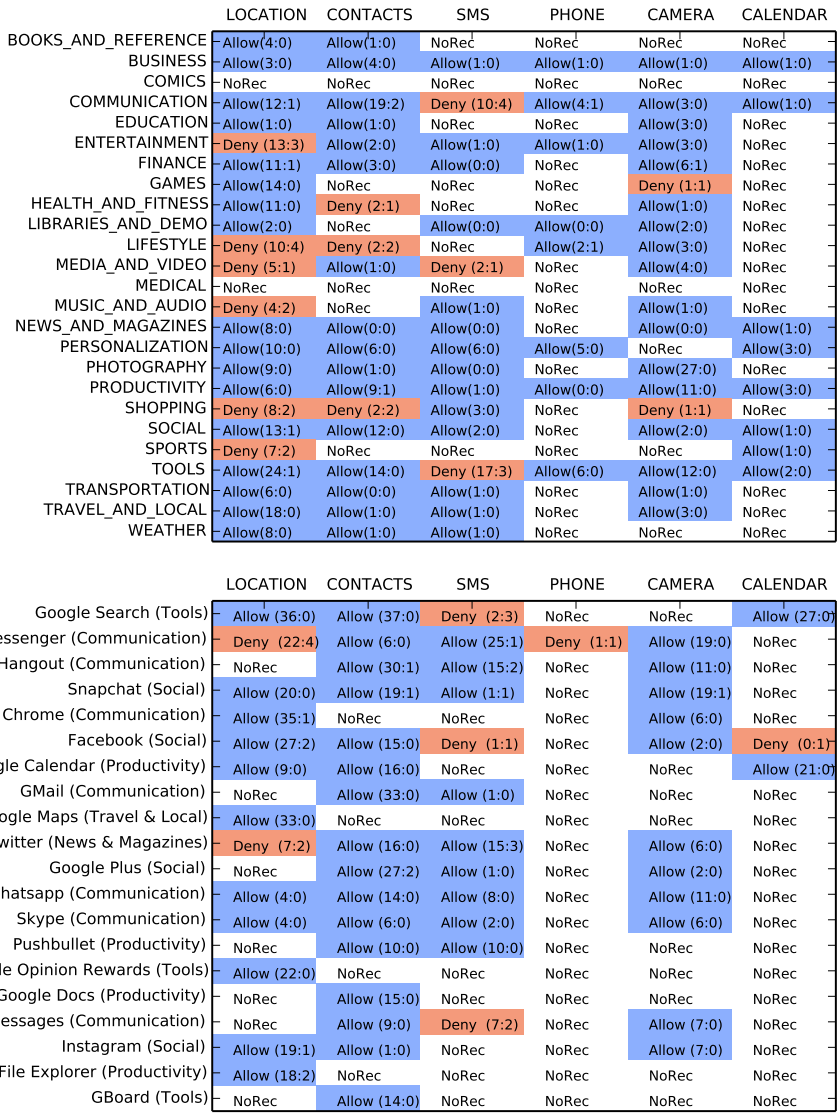

Figure 6.6: Expected Recommendations for Users in Cluster 2 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

From Figure 6.6, we find that the PPA generally recommends "Allow" for communication apps requesting access to location data. However, for one communication app, namely "Facebook Messenger," the PPA recommends a "Deny" decision.

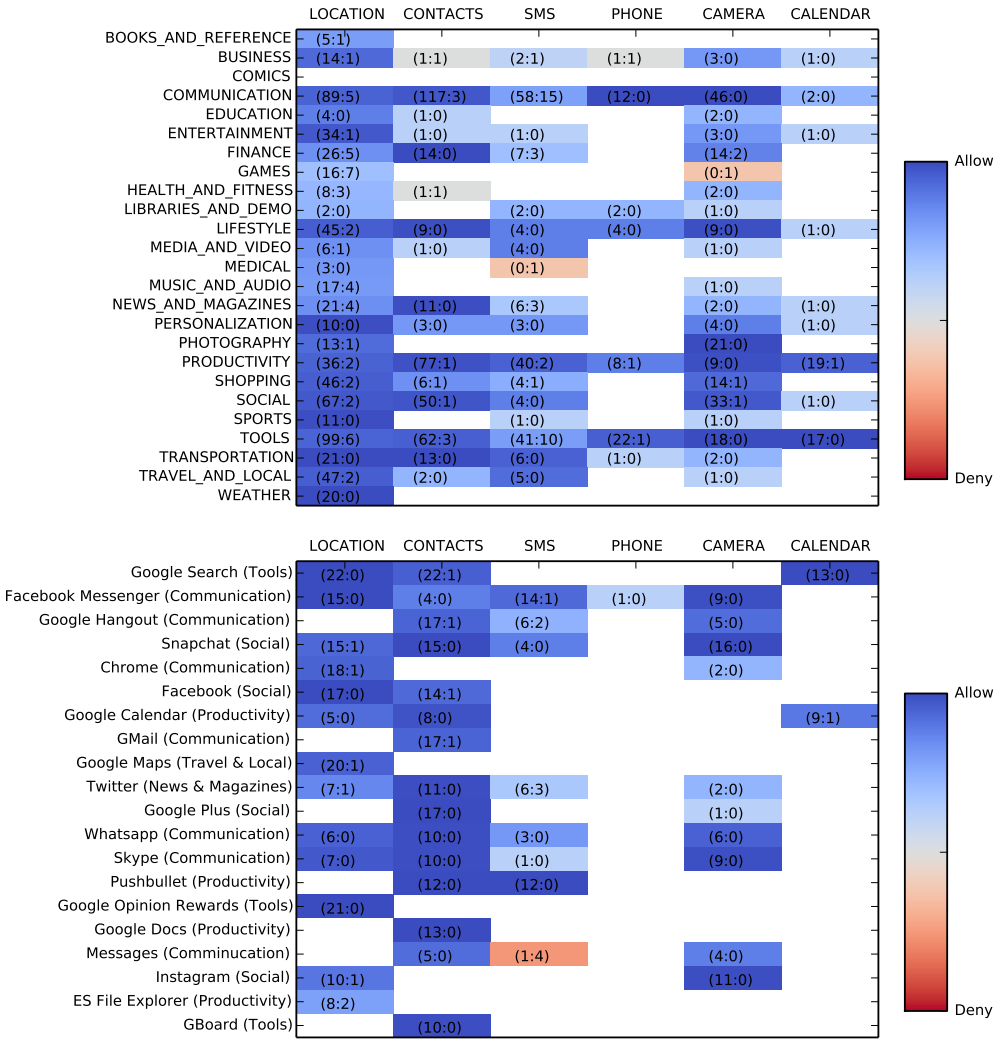

Figure 6.7: Permission Settings of Users in Cluster 3 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

This group of users (25 out of 156 users) in Figure 6.7 have relatively permissive settings for the majority of location and contact permissions, which is similar to people in Clusters 1 and 2. However, compared with Clusters 1 and 2, users in Cluster 3 are a lot less comfortable allowing game apps to access their location data, for instance.

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
BOOKS_AND_REFERENCE	Allow(4:1)	NoRec	NoRec	NoRec	NoRec	NoRec
BUSINESS	Allow(14:1)	Deny (1:1)	Deny (2:1)	Deny (1:1)	Allow(3:0)	Allow(1:0)
COMICS	NoRec	NoRec	NoRec	NoRec	NoRec	NoRec
COMMUNICATION	Allow(14:1)	Allow(23:0)	Allow(15:0)	Allow(7:0)	Allow(4:0)	Allow(0:0)
EDUCATION	Allow(4:0)	Allow(1:0)	NoRec	NoRec	Allow(2:0)	NoRec
ENTERTAINMENT	Allow(17:1)	Allow(1:0)	Allow(1:0)	NoRec	Allow(3:0)	Allow(1:0)
FINANCE	Deny (8:3)	Allow(10:0)	Allow(4:0)	NoRec	Allow(8:1)	NoRec
GAMES	Deny (16:7)	NoRec	NoRec	NoRec	Deny (0:1)	NoRec
HEALTH_AND_FITNESS	Allow(2:1)	Deny (1:1)	NoRec	NoRec	Allow(0:0)	NoRec
LIBRARIES_AND_DEMO	Allow(2:0)	NoRec	Allow(0:0)	Allow(0:0)	Allow(1:0)	NoRec
LIFESTYLE	Allow(38:2)	Allow(9:0)	Allow(4:0)	Allow(4:0)	Allow(6:0)	Allow(1:0)
MEDIA_AND_VIDEO	Allow(6:1)	Allow(1:0)	Allow(4:0)	NoRec	Allow(1:0)	NoRec
MEDICAL	Allow(3:0)	NoRec	Allow(0:1)	NoRec	NoRec	NoRec
MUSIC_AND_AUDIO	Deny (9:2)	NoRec	NoRec	NoRec	Allow(0:0)	NoRec
NEWS_AND_MAGAZINES	Allow(6:0)	Allow(0:0)	Allow(0:0)	NoRec	Allow(0:0)	Allow(1:0)
PERSONALIZATION	Allow(8:0)	Allow(0:0)	Allow(0:0)	NoRec	Allow(4:0)	Allow(1:0)
PHOTOGRAPHY	Allow(9:1)	NoRec	NoRec	NoRec	Allow(21:0)	NoRec
PRODUCTIVITY	Allow(13:0)	Allow(24:1)	Allow(10:1)	Allow(4:1)	Allow(7:0)	Allow(10:0)
SHOPPING	Allow(22:1)	Allow(5:1)	Allow(4:1)	NoRec	Allow(5:0)	NoRec
SOCIAL	Allow(15:0)	Allow(3:0)	Allow(0:0)	NoRec	Allow(4:1)	Allow(1:0)
SPORTS	Allow(11:0)	NoRec	Allow(1:0)	NoRec	Allow(1:0)	NoRec
TOOLS	Allow(19:0)	Allow(20:0)	Allow(19:3)	Allow(7:0)	Allow(11:0)	Allow(1:0)
TRANSPORTATION	Allow(11:0)	Allow(4:0)	Allow(3:0)	Allow(1:0)	Allow(0:0)	NoRec
TRAVEL_AND_LOCAL	Allow(25:1)	Allow(0:0)	Allow(3:0)	NoRec	Allow(1:0)	NoRec
WEATHER	Allow(20:0)	NoRec	NoRec	NoRec	NoRec	NoRec

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
Google Search (Tools)	Allow (22:0)	Allow (22:1)	NoRec	NoRec	NoRec	Allow (13:0)
Facebook Messenger (Communication)	Allow (15:0)	Allow (4:0)	Allow (14:1)	Allow (1:0)	Allow (9:0)	NoRec
Google Hangout (Communication)	NoRec	Allow (17:1)	Deny (6:2)	NoRec	Allow (5:0)	NoRec
Snapchat (Social)	Allow (15:1)	Allow (15:0)	Allow (4:0)	NoRec	Allow (16:0)	NoRec
Chrome (Communication)	Allow (18:1)	NoRec	NoRec	NoRec	Allow (2:0)	NoRec
Facebook (Social)	Allow (17:0)	Allow (14:1)	NoRec	NoRec	NoRec	NoRec
Google Calendar (Productivity)	Allow (5:0)	Allow (8:0)	NoRec	NoRec	NoRec	Allow (9:1)
GMail (Communication)	NoRec	Allow (17:1)	NoRec	NoRec	NoRec	NoRec
Google Maps (Travel & Local)	Allow (20:1)	NoRec	NoRec	NoRec	NoRec	NoRec
Twitter (News & Magazines)	Allow (7:1)	Allow (11:0)	Deny (6:3)	NoRec	Allow (2:0)	NoRec
Google Plus (Social)	NoRec	Allow (17:0)	NoRec	NoRec	Allow (1:0)	NoRec
Whatsapp (Communication)	Allow (6:0)	Allow (10:0)	Allow (3:0)	NoRec	Allow (6:0)	NoRec
Skype (Communication)	Allow (7:0)	Allow (10:0)	Allow (1:0)	NoRec	Allow (9:0)	NoRec
Pushbullet (Productivity)	NoRec	Allow (12:0)	Allow (12:0)	NoRec	NoRec	NoRec
Google Opinion Rewards (Tools)	Allow (21:0)	NoRec	NoRec	NoRec	NoRec	NoRec
Google Docs (Productivity)	NoRec	Allow (13:0)	NoRec	NoRec	NoRec	NoRec
Messages (Communication)	NoRec	Allow (5:0)	Deny (1:4)	NoRec	Allow (4:0)	NoRec
Instagram (Social)	Allow (10:1)	NoRec	NoRec	NoRec	Allow (11:0)	NoRec
ES File Explorer (Productivity)	Deny (8:2)	NoRec	NoRec	NoRec	NoRec	NoRec
GBoard (Tools)	NoRec	Allow (10:0)	NoRec	NoRec	NoRec	NoRec

Figure 6.8: Expected Recommendations for Users in Cluster 3 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

Figure 6.8 shows that, for users in this cluster, the PPA generally recommends allowing communication apps to access SMS. There is however a notable exception, namely "Google Hangout," for which the PPA recommends denying access to SMS. Given that the amount of data available to make these recommendations remains fairly small, it is always possible that with a larger corpus of permission decisions the PPA would reach different decisions for some of these apps.

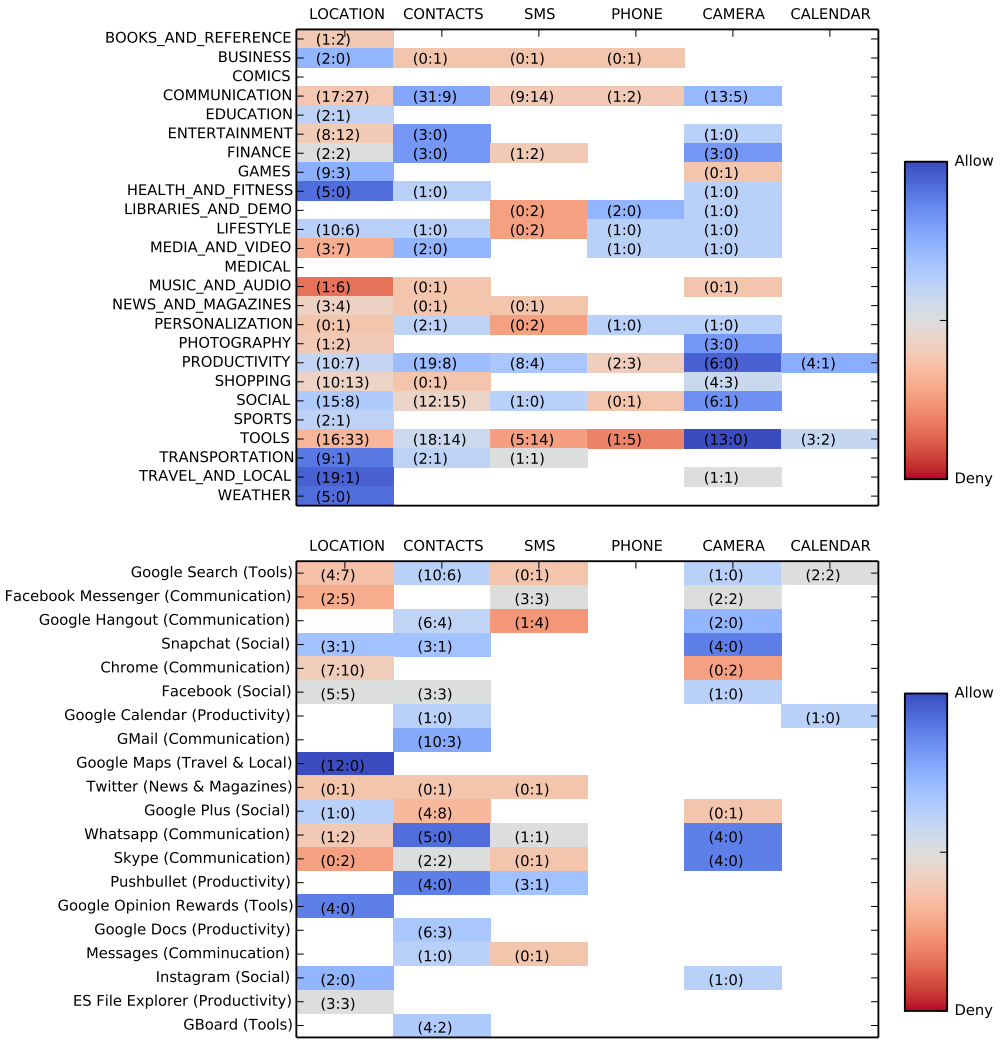

Figure 6.9: Permission Settings of Users in Cluster 4 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

According to Figure 6.9, users in this cluster (21 out of 156 users) mostly configured permission settings in a relatively conservative manner. However, even in this case, users still mostly allowed "Google Maps" to access location and "Gmail" to access contact lists. It is interesting to observe that these privacy-protective users did deny the majority of the permission settings but at the same time allowed some permission requests. From our observations, these permission requests were mostly permissions that users would expect to see, as they can reasonably be assumed to be necessary for the apps to function.

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
BOOKS_AND_REFERENCE	Deny (1:1)	NoRec	NoRec	NoRec	NoRec	NoRec
BUSINESS	Allow(2:0)	Deny (0:1)	Deny (0:1)	Deny (0:1)	NoRec	NoRec
COMICS	NoRec	NoRec	NoRec	NoRec	NoRec	NoRec
COMMUNICATION	Deny (5:5)	Allow(3:0)	Deny (2:1)	Allow(1:0)	Deny (0:1)	NoRec
EDUCATION	Deny (2:1)	NoRec	NoRec	NoRec	NoRec	NoRec
ENTERTAINMENT	Deny (7:5)	Allow(3:0)	NoRec	NoRec	Allow(1:0)	NoRec
FINANCE	Deny (0:1)	Allow(2:0)	Deny (0:2)	NoRec	Allow(1:0)	NoRec
GAMES	Deny (9:3)	NoRec	NoRec	NoRec	Deny (0:1)	NoRec
HEALTH_AND_FITNESS	Allow(3:0)	Allow(1:0)	NoRec	NoRec	Allow(1:0)	NoRec
LIBRARIES_AND_DEMO	NoRec	NoRec	Deny (0:1)	Allow(1:0)	Allow(1:0)	NoRec
LIFESTYLE	Deny (9:6)	Allow(1:0)	Deny (0:2)	Allow(1:0)	Allow(1:0)	NoRec
MEDIA_AND_VIDEO	Deny (3:7)	Allow(2:0)	NoRec	Deny (1:0)	Allow(1:0)	NoRec
MEDICAL	NoRec	NoRec	NoRec	NoRec	NoRec	NoRec
MUSIC_AND_AUDIO	Deny (1:4)	Deny (0:1)	NoRec	NoRec	Deny (0:1)	NoRec
NEWS_AND_MAGAZINES	Deny (1:1)	Deny (0:0)	Deny (0:0)	NoRec	NoRec	NoRec
PERSONALIZATION	Deny (0:1)	Deny (2:1)	Deny (0:2)	Allow(1:0)	Allow(1:0)	NoRec
PHOTOGRAPHY	Deny (1:2)	NoRec	NoRec	NoRec	Allow(3:0)	NoRec
PRODUCTIVITY	Deny (4:3)	Allow(4:1)	Deny (2:1)	Deny (2:2)	Allow(6:0)	Allow(3:0)
SHOPPING	Deny (6:6)	Deny (0:1)	NoRec	NoRec	Allow(4:1)	NoRec
SOCIAL	Deny (3:2)	Deny (2:3)	Allow(1:0)	Deny (0:1)	Allow(0:0)	NoRec
SPORTS	Deny (2:1)	NoRec	NoRec	NoRec	NoRec	NoRec
TOOLS	Deny (3:12)	Deny (3:2)	Deny (3:4)	Deny (0:2)	Allow(7:0)	Allow(1:0)
TRANSPORTATION	Allow(7:0)	Allow(0:0)	Deny (0:1)	NoRec	NoRec	NoRec
TRAVEL_AND_LOCAL	Deny (6:1)	NoRec	NoRec	NoRec	Deny (1:1)	NoRec
WEATHER	Allow(5:0)	NoRec	NoRec	NoRec	NoRec	NoRec

	LOCATION	CONTACTS	SMS	PHONE	CAMERA	CALENDAR
Google Search (Tools)	Deny (4:7)	Deny (10:6)	Deny (0:1)	NoRec	Allow (1:0)	Deny (2:2)
Facebook Messenger (Communication)	Deny (2:5)	NoRec	Deny (3:3)	NoRec	Deny (2:2)	NoRec
Google Hangout (Communication)	NoRec	Deny (6:4)	Deny (1:4)	NoRec	Allow (2:0)	NoRec
Snapchat (Social)	Deny (3:1)	Deny (3:1)	NoRec	NoRec	Allow (4:0)	NoRec
Chrome (Communication)	Deny (7:10)	NoRec	NoRec	NoRec	Deny (0:2)	NoRec
Facebook (Social)	Deny (5:5)	Deny (3:3)	NoRec	NoRec	Allow (1:0)	NoRec
Google Calendar (Productivity)	NoRec	Allow (1:0)	NoRec	NoRec	NoRec	Allow (1:0)
GMail (Communication)	NoRec	Allow (10:3)	NoRec	NoRec	NoRec	NoRec
Google Maps (Travel & Local)	Allow (12:0)	NoRec	NoRec	NoRec	NoRec	NoRec
Twitter (News & Magazines)	Deny (0:1)	Deny (0:1)	Deny (0:1)	NoRec	NoRec	NoRec
Google Plus (Social)	Allow (1:0)	Deny (4:8)	NoRec	NoRec	Deny (0:1)	NoRec
Whatsapp (Communication)	Deny (1:2)	Allow (5:0)	Deny (1:1)	NoRec	Allow (4:0)	NoRec
Skype (Communication)	Deny (0:2)	Deny (2:2)	Deny (0:1)	NoRec	Allow (4:0)	NoRec
Pushbullet (Productivity)	NoRec	Allow (4:0)	Deny (3:1)	NoRec	NoRec	NoRec
Google Opinion Rewards (Tools)	Allow (4:0)	NoRec	NoRec	NoRec	NoRec	NoRec
Google Docs (Productivity)	NoRec	Deny (6:3)	NoRec	NoRec	NoRec	NoRec
Messages (Communication)	NoRec	Allow (1:0)	Deny (0:1)	NoRec	NoRec	NoRec
Instagram (Social)	Allow (2:0)	NoRec	NoRec	NoRec	Allow (1:0)	NoRec
ES File Explorer (Productivity)	Deny (3:3)	NoRec	NoRec	NoRec	NoRec	NoRec
GBoard (Tools)	NoRec	Deny (4:2)	NoRec	NoRec	NoRec	NoRec

Figure 6.10: Expected Recommendations for Users in Cluster 4 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

From Figure 6.10, we can see that the PPA generally recommends "Deny" if travel & local apps request to access location data. However, for the travel & local app "Google Maps," the PPA provides an "Allow" recommendation

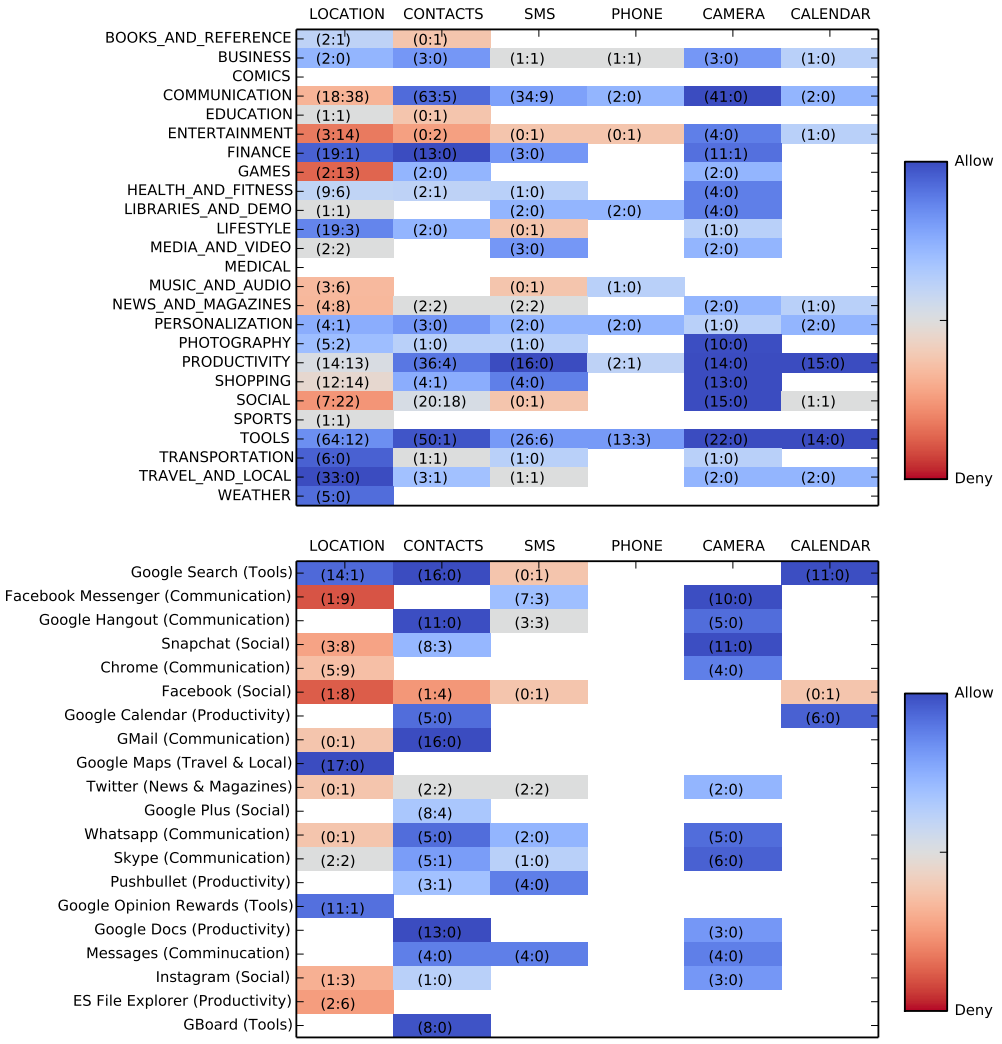

Figure 6.11: Permission Settings of Users in Cluster 5 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

Figure 6.11 shows that the users in this cluster (17 out of 156 users) had most of their protective settings on location data access. These users also had relatively conservative settings on social apps (such as "Facebook," "Google+," and "Snapchat") accessing contact lists, compared with other profiles. Different from the most conservative Cluster 4, users in Cluster 5 generally allow camera and calendar permissions.

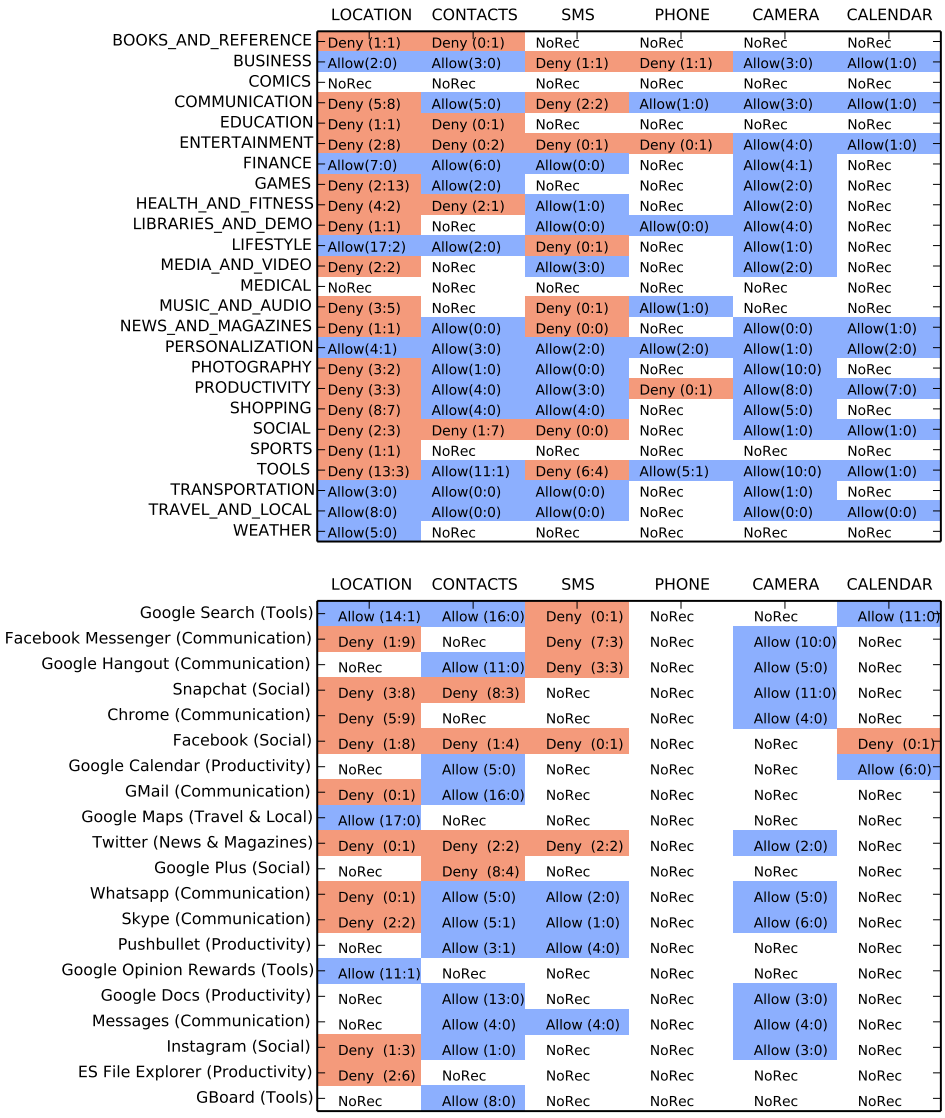

Figure 6.12: Expected Recommendations for Users in Cluster 5 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

From Figure 6.12, we see that for this cluster, the PPA generally recommends allowing social apps to access calendar data. However, for one social app, namely "Facebook," the PPA actually recommends denying access. In this particular case however this recommendation is effectively based on a single data point. This is a situation where it might actually be better to just prompt the user rather than offer a recommendation. This could be done using a threshold below which if the PPA does not have enough data points, it would refrain from making a recommendation and would just prompt the user for a decision.

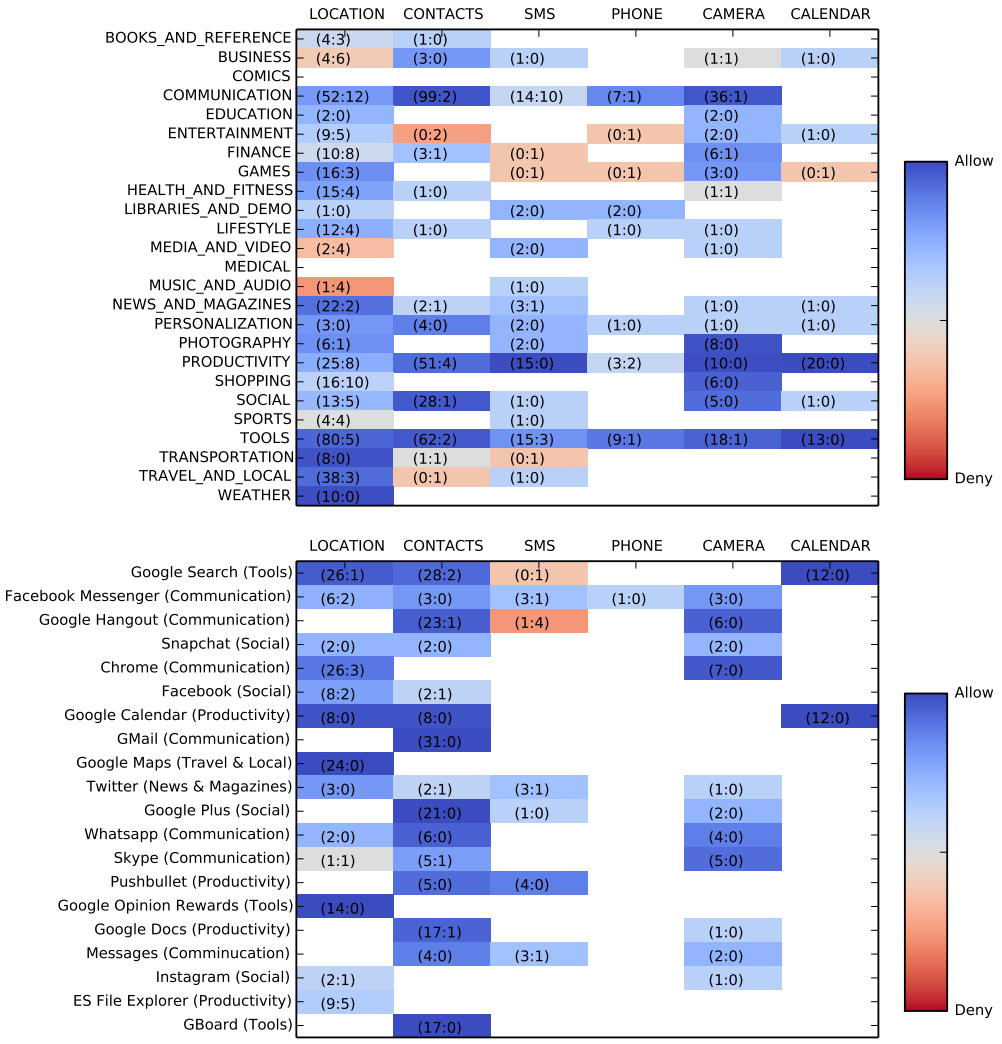

Figure 6.13: Permission Settings of Users in Cluster 6 (App Categories and Top 20 Apps)
We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. The color of each cell indicates overall preferences: blue for more allowing than denying; red for more denying than allowing; white for no data observed. The color depth indicates the agreement: the darker the color, the higher confidence we have from the observed data.

As shown in Figure 6.13, users in this cluster (37 out of 156 users) tend to have fairly permissive settings. However, compared with the permissive clusters we discussed (Clusters 1, 2, and 3), users in this group had more restrictions on sharing their location data. For example, we found that five out of eight of users who had the shopping app "eBay" installed denied its location access. On the other hand, compared with the most protective profile (Cluster 3), users in this group mostly allowed non-location permission access except for some anecdotal examples such as the "MyFitnessPal" app accessing the camera data.

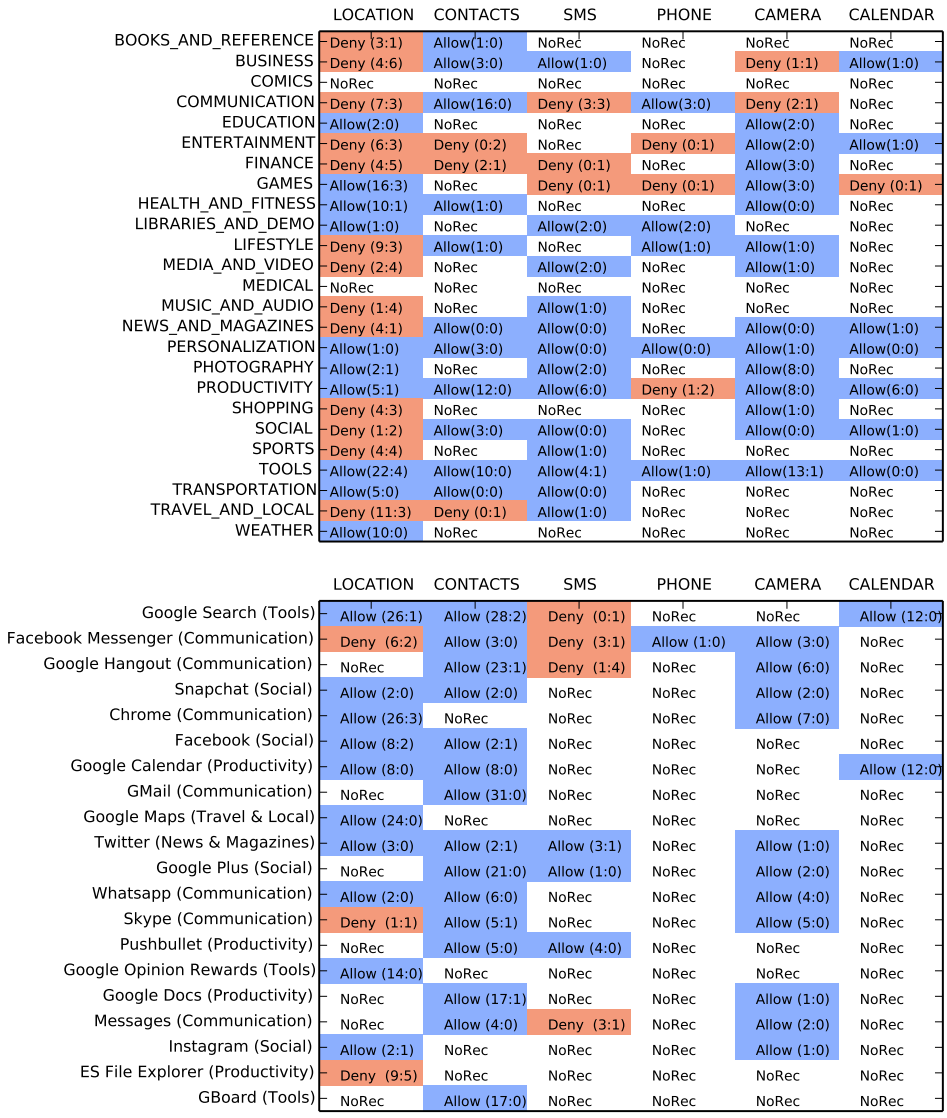

Figure 6.14: Expected Recommendations for Users in Cluster 6 (App Categories and Top 20 Apps)

We show the number of "(allow:deny)" decisions made by users in this cluster in each cell. Note that in the table above, for categories, we did not count the settings from the most popular 100 apps. The color of each cell indicates the recommendation decision: blue for allow, red for deny, and white for no recommendation because of low confidence from the classifier.

From Figure 6.14, we see that the PPA generally recommends "Deny" for communication apps requesting access to location data. However, for social app "Snapchat," the PPA recommends "Allow." As with previous clusters, some of the PPA recommendations shown in the table rely on fairly small numbers of recorded decisions. One could imagine requiring a minimum number of decisions in a given entry to rely on the PPA's recommendation and otherwise simply prompting the user for a decision.

6.4 Stats and Findings From Our App Deployment

We published the enhanced version of the PPA app described above in the Google Play store in early 2017. In its current form, the PPA app is mainly intended for Android 5.X users with rooted Android phones, as users of more recent versions of Android would need to go through several non-trivial configuration steps to get the app to run on their phones. As a result the PPA app has only been downloaded by a limited number of users – 151 verified downloads as of the time of writing, though it has a rating of 4.3. Until recently, people who downloaded the PPA had the option of completing an IRB consent form, which allowed us to collect anonymous app permission setting data. Among the 151 people who downloaded the PPA, 18 users opted to share their permission settings with us.

Among the users who shared data with us, we found that the PPA app got an overall F-1 score of 71.7% (0.76 for allow, 0.65 for deny) when predicting the permission settings for these users. Among all the deny recommendations, users manually acted on 68% of them (299 out of 442). The details of how each user interacted with the recommendations can be found in Table 6.1.

In interpreting these numbers, it is important to keep in mind that this version of the PPA app does not require users to review the recommendations made by the PPA or their current permission settings. Instead, users have to browse permission settings in the PPA app and manually toggle permission settings to accept any PPA recommendation. As a result, it is not surprising that, among the 18 users, only 10 engaged with their settings within the first three days of installing the PPA. Given the above, including the limited number of users for whom we were able to collect data, it is hard to reach additional conclusions. The results do not seem inconsistent with the ones reported earlier in this dissertation, and seem to confirm that, in general, a majority of the recommendations made by the PPA tend to be accepted by users.

We generated six clusters of users using permission settings data collected from 154 participants in our field studies, as already reported earlier. The proportion of users in each cluster was similar to that reported earlier: 17 (10.9%) in Cluster 1, 39 (25.0%) in Cluster 2, 25 (16.0%) in Cluster 3, 21 (13.5%) in Cluster 4, 17 (10.9%) in Cluster 5, and 37 (23.7%) in Cluster 6. However, among the 18 users of the PPA in the Google Play store for whom we were able to collect data, we found that half were assigned to Cluster 4, which is the most protective of the six profiles. While these are small numbers and it is hard to reach any strong conclusion based on such a small sample, this might be attributed to selection bias, namely people who are more likely to download the app and provide IRB consent might also be people who are particularly concerned about their privacy. These users might actually have found our PPA app after reading articles in the press about it (e.g., [71]) or may have manually searched for Android permission manager tools.

By analyzing the permission settings and users' reactions to the recommendations, we found that it would not be an instant process for users to configure their app permission settings prop-

Table 6.1: Statistics on Permission Settings and Recommendations
(Note: "#" is short for "number of" and "Rec" is short for "Recommended to")

User	#Apps[a]	#Settings[b]	Profile[c]	User Allow		User Deny	
				Rec Allow	Rec Deny	Rec Allow	Rec Deny
U-1	15	24	#4	9	14	0	1
U-2	33	57	#4	29	22	0	6
U-3	34	59	#2	32	17	0	10
U-4	20	29	#5	26	3	0	0
U-5	15	24	#4	7	13	0	4
U-6	47	189	#4	42	2	70	77
U-7	21	36	#3	26	0	10	0
U-8	25	73	#2	18	10	24	21
U-9	22	42	#4	18	3	10	12
U-10	67	111	#4	55	32	7	17
U-11	42	61	#4	25	1	0	35
U-12	75	107	#4	65	5	3	34
U-13	29	47	#2	43	1	0	3
U-14	40	69	#4	15	2	19	36
U-15	25	73	#2	18	10	24	21
U-16	28	50	#3	42	2	6	0
U-17	22	48	#2	17	3	10	18
U-18	15	24	#2	17	3	0	4

[a] We did not include system apps (apps that have system signatures and have been pre-installed on users' devices) or non-Google-Play apps in our analysis. And the PPA app manages six permissions: location, contacts, SMS, phone, camera, and calendar.

[b] We did not use privacy nudges to motivate users to engage with their settings or give users a deadline to review their settings.

[c] This number shows the dominating profile assigned to this user. By analyzing their answers to the profile-assignment questionnaire, the PPA app generates a weighted vector representing the profile assignment. Here we report only the dominant assignment. (Cluster #2 is the most permissive; Cluster #4 is the most protective.)

erly. Users could be actively installing and trying new apps, playing around with different permission settings, or changing the settings as their privacy preferences change over time (for reasons such as gaining experience using Android, or being educated by privacy incidents or information about privacy protection). Thus, it is crucial to design the privacy assistant tools to keep interacting with users in the long term, to capture the changes needed in privacy settings in an agile way.

Chapter 7

Conclusions

7.1 Summary of Contribution

The focus of this dissertation was to determine to what extent machine learning techniques can help users manage their mobile app privacy settings. This included looking at whether machine learning techniques can be used to predict people's mobile app permission settings. As part of our work, we have quantified the predictive power of different machine learning techniques. We have also explored how machine learning models can possibly help reduce the otherwise unrealistically high burden associated with the large number of mobile app permissions people are expected to configure. As part of this research, we have evaluated initial configurations of machine learning functionality that rely on recommendations as a way of providing people with the benefits of machine learning's predictive power without taking away control from users.

The main contributions of this work are summarized below.

- We demonstrated that it is possible to build models of people's privacy preferences that can help predict many mobile app permission settings. Using a large dataset of Android permission settings from 4.8 million Android phone users using LBE Privacy Guard, we showed that, while people's privacy preferences are complex and diverse, these preferences do lend themselves to the development of machine learning models with strong predictive power. As part of a first study where we focused on this large dataset, we were able to show that machine learning can in theory help predict many of a user's permission preferences and, as a result, can also help reduce user burden when it comes to configuring mobile app permission settings. This included simulating functionality that would use predictions from the machine learning models when we have sufficient confidence in these predictions, and falling back on asking users when we do not have sufficient confidence in the model's predictions. This simulation study suggested that, at least in theory, it would be possible to achieve accuracy of well over 90 percent, while drastically reducing the number of privacy

decisions users have to make.

- When limiting ourselves to simple clustering techniques, we found that while users' privacy preferences are diverse, a small number of clusters can go a long way in capturing and predicting people's mobile app permission settings.
- Because our first study was conducted using data obtained from users who had not been explicitly encouraged to interact with their settings, we felt that we should confirm the results of our first study using permission settings collected from users who would be explicitly nudged to review and possibly modify their permission settings. While conducting such a study could only be done with a significantly smaller number of users, the data collected through this second study confirmed our initial findings. Here again, we were able to show that while people's privacy preferences are diverse, a small number of clusters can go a long way in helping predict their mobile app permission settings.
- Beyond showing that people's mobile app permission settings can often be predicted, this dissertation went one step further and actually developed and tested a privacy assistant designed to leverage our findings. Part of the challenge in developing such an assistant involved determining how to best leverage the predictive power of machine learning without taking control away from users. We adopted a simple approach, where predictions made by the privacy assistant are turned into recommendations that the user can review and edit prior to optionally accepting them. The assistant relied on a simple set of clusters. The assistant asks the user a small number of questions to assign him or her to a cluster. Each cluster is itself associated with a set of mobile app permission recommendations, organized by categories of apps and types of permissions. These recommendations are instantiated based on the actual apps a user has on his or her smartphone. Evaluation of the resulting privacy assistant showed that users accepted the majority of the recommendations made by their privacy assistant; they felt comfortable continuing to operate with the resulting permission settings and were generally quite pleased with the privacy assistant functionality, with users reporting that they appreciated the reduction in user burden.

  We used the profiles built by our privacy assistant and analyzed how users in different clusters differ from one another. This analysis shows that while some users are very permissive and others are very conservative, a number of people have more nuanced privacy preferences. Despite these differences, people within each cluster tend to agree on many privacy decisions. This commonality is the basis for the predictive power of the privacy assistants we piloted in our study. In the future, one could envision more complex privacy assistants that learn from their interactions with their users, similar to some of the more personalized models with which we experimented using our initial corpus of LBE users.

  Finally, this dissertation also resulted in a refined version of our mobile app privacy assistant being developed and released on the Google Play store. While this app is only available to users of rooted Android phone and, as a result, has only had a limited number

of downloads, analysis of data collected from the small number of users who provided IRB consent, seems to further confirm the usefulness of our mobile app privacy assistant.

7.2 Limitations of This Work

This dissertation involved data collected from three different sets of Android users: LBE Privacy Guard users, users of rooted Android phones who were exposed to privacy nudges, and users of our personalized privacy assistants, who were also all using rooted Android phones. LBE Privacy Guard users were mainly users who used phones with the MIUI Android ROM or used rooted Android phones, with many of these users being based in China. Data collected as part of our two field studies all involved users of rooted Android phones. As such we are not in a position where we can categorically claim that the results of our study necessarily apply to the broader population of Android phone users, let alone iOS users. There is a reasonable chance that users of rooted Android phones have somewhat different privacy preferences than the rest of the Android user population (e.g., they are likely to be more technically savvy). We believe however that, while the actual privacy profiles we built for our study might have been slightly different if one were to look at the general population of Android users, there is a very good chance that these profiles would lend themselves to the identification of somewhat similar clusters. After all, while the proportion of people concerned about different types of privacy issues might be different in the general population of Android users, their privacy concerns are likely to exhibit somewhat similar correlations, which in turn would lend itself to the identification of clusters, albeit possibly somewhat different ones.

In addition to the above, our recruitment processes themselves may have introduced some biases too. We recruited participants in a voluntary way instead of choosing sample users randomly. Thus, we might have introduced selection bias in the process. Subjects recruited for our studies might have been self-motivated to participate. Thus, the theory of probability sampling might not be applied[44].

In this dissertation, we experimented with multiple machine learning models to predict users' app permission settings. Clearly, the work reported herein was limited to relatively simple techniques. Better results could possibly be obtained with more sophisticated techniques, with larger datasets, and also with more sophisticated ways of configuring interactions between privacy assistants and their users. In addition, it should be noted that the ground truth assumed to evaluate the predictive power of our models is itself imperfect. Specifically, the ground truth data we used for training and evaluation comes from users' actual settings on their phones. In our first study, we only used data from a subset of about 200,000 LBE users who seemed to be truly engaged with their settings. There is no guarantee however that the settings collected from these users fully reflect their actual privacy preferences. The same is true for data collected in our later studies. While nudging has been shown to motivate people to review and adjust their permission

settings, there is no guarantee that the settings we collected from these users fully capture their privacy preferences.

Finally, our evaluation of privacy assistants was conducted over a relatively short period of time. The research reported herein does not look at how people's privacy preferences might change over time and how privacy assistants may want to regularly check with users to see whether such changes might require updating the models on which the assistants rely for their recommendations.

7.3 Open Questions and Future Directions

Could we improve the predictive power of our models?

In this dissertation, we studied the LBE privacy guard user data at scale and conducted two field studies to pilot our profile-based privacy assistant for mobile app permissions. Our prediction model generated recommendations of allowing or denying access by a category of apps or an individual app to a specific permission. It goes without saying that the models presented in this dissertation would benefit from being trained on significantly larger corpora of data such as the corpora a company like Google or Samsung likely has access to.

A related, yet different question has to do with whether the predictive power of our clusters (or other machine learning techniques) could be improved by introducing additional features such as context information. Work in Contextual Integrity [69] as well as work quantifying tradeoffs between accuracy and user burden or work looking at the impact of permission purpose on people's privacy preferences suggests that this might be an avenue worth exploring (e.g., [21, 57, 59]). If such models were to prove to have stronger predictive power, this could also argue for the introduction of more expressive permission settings, which would allow users to distinguish between additional contextual attributes when it comes to granting different permissions to different apps (e.g., based on the purpose for which a permission is requested). Today such settings are unavailable and hence the development of such models is of little practical use. Instead, users today are left to infer on their own the potential purpose(s) for which permissions are being requested.

We acknowledge that in iOS, developers can provide short motivations to explain why they need a specific permission. Android also gives developers a signal to show custom-defined screens to explain the permission request before popping up system dialog for users to make decisions. Unfortunately, because the this information is provided in the form of free text, developers tend to fall back on rather vague statements [88]. In short, today these explanations increase complexity and user burden of app privacy management but do not give users more convenient control.

What are the best ways of exposing this functionality to users?

On the one hand, we need to think about the tradeoff between automation and user autonomy. We have shown that it is possible to build machine learning models to predict users' app permission settings accurately. One may apply the prediction results in an automatic way, such as directly changing the settings for the users, or in a passive way that still requires users' participation.

We believe that full automation is not the right way to go. We still need to make sure that users are in control. If we apply fully automated solutions to configure the permission settings directly, it would be difficult for users to participate in the process and retain their agency. In the field studies of this dissertation, we applied the predictions of permission settings as recommendations for users to review. Users can decide whether to accept or reject our recommendations. And users can also change their decisions later. Future research should explore different possible ways of combining automation and user control, through different types of user dialogues and different types of interfaces, looking at accuracy, user burden, and overall sense of control by users. Intelligent tools need users' input, such as privacy settings, attitudes, or reactions to data collection, to infer their preferences and provide assistance. Collecting more data from users may result in a significant increase in user burden. One way to better utilize users' engagement would be by asking for user input about the most informative or most confusing decisions. In Chapter 3, we simulated scenarios where the system prompts users to give feedback on permission settings that the model has a hard time predicting with sufficient confidence. It would be interesting to further experiment with different strategies and scenarios to find a good balance of volume of input data for prediction models and user burden.

Can we extend this methodology to other domains?

In this dissertation, we chose mobile app permissions as the domain to study ways to reconcile usability and privacy. The nature of high user burden can be found in other domains, such as content-sharing control of social networks, website access settings of browsers, and access control of Internet of Things (IoT) devices.

We naturally have two open questions about extending our technology to other domains:

- Can we apply our methodology to obtain similar results in other domains?

  Social networks such as Facebook have rich collections of settings available to users about whether to share a specific post with their friends, friends of friends, or everyone on the network. IoT devices, such as smart speakers, also provide options for users to control data access by various third parties. For example, a speaker may need email or calendar access if a user needs a restaurant booking service enabled.

  We expect these scenarios have a similar nature compared with mobile app permissions in terms of personalized preferences and predictive modeling. It would be interesting to study how our methodology of privacy profiles and profile-based recommendations can assist users in these domains.

- Do privacy profiles built for mobile app permissions extend to other domains?

  Open a brand new box, start the system, and configure all the settings one by one from scratch: this is a cold-start process for most of us when we start using a new device. In another scenario, where a user wants to make certain changes to privacy settings, the user would have to adjust the configurations on all of their many smart devices and services one by one. Thus, a user might be alerted by an uncomfortable location exposure and turn off location access on the smartphone, but still be sharing location with fitness trackers if the user forgot to take care of them.

  Would it possible for us to model users' diverse privacy preferences from one domain, and use that knowledge to give a warm-start when configuring settings for other domains as well? For example, can we learn patterns so that we can suggest the user to turn off facial recognition on a thermometer, if the user denied camera access of apps on their smartphone? Or can we learn an explicit or hidden feature representation (such as clusters of users) that can describe users' diverse privacy preferences so that we can directly apply them to new services?

List of Figures

List of Tables

Bibliography

[1] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, 1980. URL http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm#guidelines.

[2] Health Insurance Portability and Accountability Act of 1996 (HIPAA), 1996. URL https://www.govinfo.gov/content/pkg/PLAW-104publ191/pdf/PLAW-104publ191.pdf.

[3] Children's Online Privacy Protection Rule (COPPA), 1998. URL https://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule.

[4] The Online Privacy Protection Act of 2003, Cal. Bus. & Prof. Code §§ 22575-22579, 2003. URL https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?lawCode=BPC&sectionNum=22575.

[5] Fair Information Practice Principles (Archived), 2007. URL https://web.archive.org/web/20090331134113/http://www.ftc.gov/reports/privacy3/fairinfo.shtm.

[6] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Da, 2016. URL https://eur-lex.europa.eu/eli/reg/2016/679/oj.

[7] Evrim Acar, Daniel Dunlavy, Tamara Kolda, and Morten Mørup. Scalable Tensor Factorizations with Missing Data. In *Proc. of 10th SIAM Int. Conf. on Data Mining*, pages 701–712, 2010. doi: 10.1137/1.9781611972801.61.

[8] Mark S. Ackerman, Lorrie Faith Cranor, and Joseph Reagle. Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences. *Proceedings of the 1st ACM conference on Electronic Commerce - EC '99*, pages 1–8, 1999. ISSN 1581131763. doi: 10.1145/336992.336995. URL http://dl.acm.org/citation.cfm?id=336992.336995.

[9] Alessandro Acquisti and Jens Grossklags. Privacy and Rationality in Individual Decision Making. *IEEE Security and Privacy*, 3(1):26–33, 1 2005. ISSN 1540-7993. doi: 10.1109/MSP.2005.22. URL http://dx.doi.org/10.1109/MSP.2005.22.

[10] Alessandro Acquisti, Idris Adjerid, Rebecca Hunt Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM CSUR*, 1(40), 2016.

[11] Idris Adjerid, Eyal Peer, and Alessandro Acquisti. Beyond the Privacy Paradox: Objective versus Relative Risk in Privacy Decision Making. *SSRN*, 2016. doi: 10.2139/ssrn.2765097. URL https://ssrn.com/abstract=2765097.

[12] Yuvraj Agarwal and Malcolm Hall. ProtectMyPrivacy: Detecting and Mitigating Privacy Leaks on iOS Devices Using C rowdsourcing. *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services - MobiSys '13*, 6 (September):97, 2013. doi: 10.1145/2462456.2464460. URL http://dl.acm.org/citation.cfm?doid=2462456.2464460.

[13] Hazim Almuhimedi. *Helping Smartphone Users Manage their Privacy through Nudges*. PhD thesis, 2017.

[14] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. *Proc. of the 2015 ACM Conference on Human Factors in Computing Systems (CHI)*, pages 787–796, 2015. doi: 10.1145/2702123.2702210.

[15] App Annie. The Average Smartphone User Accessed Close to 40 Apps per Month in 2017, 2017. URL https://www.appannie.com/en/insights/market-data/apps-used-2017/.

[16] Michelle Atkinson. Apps Permissions in the Google Play Store — Pew Research Center, 2015. URL http://www.pewinternet.org/2015/11/10/apps-permissions-in-the-google-play-store/.

[17] Gkhan Bal, Kai Rannenberg, and Jason I. Hong. Styx: Privacy Risk Communication for the Android Smartphone Platform Based on Apps' Data-Access Behavior Patterns. *Computers and Security*, 53(69):187–202, 2015. ISSN 01674048. doi: 10.1016/j.cose.2015.04.004.

[18] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. "Little Brothers Watching You": Raising Awareness of Data Leaks on Smartphones. *SOUPS '13: Proceedings of the Ninth Symposium on Usable Privacy and Security*, pages 12:1–12:11, 2013. doi: 10.1145/2501604.2501616. URL http://doi.acm.org/10.1145/2501604.2501616.

[19] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. The Impact of Timing on the Salience of Smartphone App Privacy Notices. *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices - SPSM '15*, pages 63–74, 2015. doi: 10.1145/2808117.2808119. URL http:

//dl.acm.org/citation.cfm?id=2808117.2808119.

[20] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. Privacy and Contextual Integrity: Framework and Applications. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, SP '06, pages 184–198, Washington, DC, USA, 2006. IEEE Computer Society. ISBN 0-7695-2574-1. doi: 10.1109/SP.2006.32. URL http://dx.doi.org/10.1109/SP.2006.32.

[21] Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, 2011. ISSN 16174909. doi: 10.1007/s00779-010-0346-0.

[22] Alastair R Beresford, Andrew Rice, and Nicholas Skehin. MockDroid: Trading Privacy for Application Functionality on Smartphones Categories and Subject Descriptors. *HotMobile '11 Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*, pages 49–54, 2011. doi: 10.1145/2184489.2184500.

[23] Tom Buchanan, Carina Paine, Adam N Joinson, and Ulf-Dietrich Reips. Development of Measures of Online Privacy Concern and Protection for Use on the Internet. *J. Am. Soc. Inf. Sci. Technol.*, 58(2):157–165, 1 2007. ISSN 1532-2882. doi: 10.1002/asi.v58:2. URL http://dx.doi.org/10.1002/asi.v58:2.

[24] Ann Cavoukian. Privacy by Design: The 7 Foundational Principles, 2009. URL https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf.

[25] Farah Chanchary and Sonia Chiasson. User Perceptions of Sharing, Advertising, and Tracking. *Symposium on Usable Privacy and Security (SOUPS) 2015, July 2224*, pages 53–67, 2015.

[26] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. Nudging People Away from Privacy-Invasive Mobile Apps through Visual Framing. *INTERACT*, III(LNCS 8119):74–91, 2013.

[27] Comscore. Comscore Reports January 2016 U.S. Smartphone Subscriber Market Share, 2016. URL https://www.comscore.com/Insights/Rankings/comScore-Reports-January-2016-US-Smartphone-Subscriber-Market-Share.

[28] Sunny Consolvo, Ian E Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. Location Disclosure to Social Relations: Why, When, & What People Want to Share. *CHI 2005 Conference on Human Factors in Computing Systems*, pages 81–90, 2005. ISSN 02749696. doi: 10.1145/1054972.1054985.

[29] Lorrie Faith Cranor and Joseph Reagle. Beyond Concern: Understanding Net Users' Attitudes About Online Privacy. Technical Report August, 1999.

[30] Justin Cranshaw, Jonathan Mugan, and Norman Sadeh. User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models. *AAAI*, pages 1146–1152, 2011.

URL http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3785.

[31] Janna Lynn Dupree, Richard Devries, Daniel M Berry, and Edward Lank. Privacy Personas: Clustering Users via Attitudes and Behaviors toward Security Practices. *Conference on Human Factors in Computing Systems*, pages 5228–5239, 2016. doi: 10.1145/2858036.2858214.

[32] Serge Egelman and Eyal Peer. Predicting Privacy and Security Attitudes. *Computers and Society: The Newletter of ACM SIGCAS*, 45(1):22–28, 2015. ISSN 00952737. doi: 10.1145/2738210.2738215. URL https://www.icsi.berkeley.edu/icsi/publication_details?n=3738.

[33] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones. *Communications of the ACM*, 57(3):99–106, 2014. ISSN 00010782. doi: 10.1145/2494522. URL http://dl.acm.org/citation.cfm?doid=2566590.2494522.

[34] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. LIBLINEAR: A Library for Large Linear Classification. *Journal of Machine Learning Research*, 9(2008):1871–1874, 2008. ISSN 15324435. doi: 10.1038/oby.2011.351.

[35] Lujun Fang and Kristen LeFevre. Privacy Wizards for Social Networking Sites. *Proceedings of the 19th International Conference on World Wide Web*, page 351, 2010. ISSN 15360695. doi: 10.1145/1772690.1772727. URL http://dl.acm.org.eaccess.ub.tum.de/citation.cfm?id=1772690.1772727.

[36] Adrienne Porter Felt, Serge Egelman, and David Wagner. I've Got 99 Problems, But Vibration Ain't One: A Survey of Smartphone Users' Concerns. *Proceedings of the 2nd ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 33–44, 2012. ISSN 15437221. doi: 10.1145/2381934.2381943.

[37] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android Permissions: User Attention, Comprehension, and Behavior. *SOUPS '12: Proceedings of the Eighth Symposium on Usable Privacy and Security*, pages 1–14, 2012. ISSN 09581669. doi: http://doi.acm.org/10.1145/2335356.2335360.

[38] Huiqing Fu, Yulong Yang, Nileema Shingte, Janne Lindqvist, and Marco Gruteser. A Field Study of Run-Time Location Access Disclosures on Android Smartphones. *USEC*, (February), 2014. doi: 10.14722/usec.2014.23044. URL http://www.winlab.rutgers.edu/~janne/USECfieldstudy.pdf.

[39] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2647–2656, 2014. doi: 10.1145/2556288.2556978. URL http://dl.acm.org/citation.cfm?

id=2556288.2556978.

[40] Benjamin Henne, Christian Kater, and Matthew Smith. On Usable Location Privacy for Android with Crowd-Recommendations. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8564 LNCS:74–82, 2014. ISSN 16113349. doi: 10.1007/978-3-319-08593-7{_}5.

[41] Chris Jay Hoofnagle and Jennifer M. Urban. *Alan Westin's Privacy Homo Economicus*, volume 49. 2014. ISBN 1999082702. URL http://scholarship.law.berkeley.edu/facpubs/2395.

[42] IAPP. Fair Information Practice Principles. URL https://iapp.org/resources/article/fair-information-practices/.

[43] Qatrunnada Ismail, Tousif Ahmed, Apu Kapadia, and Michael K Reiter. Crowdsourced Exploration of Security Configurations. *Proceedings of the ACM CHI'15 Conference on Human Factors in Computing Systems*, 1:467–476, 2015. doi: 10.1145/2702123.2702370. URL http://dx.doi.org/10.1145/2702123.2702370.

[44] Bethlehem Jelke. Selection Bias in Web Surveys. *International Statistical Review*, 78 (2):161–188, 2010. ISSN 0306-7734. doi: 10.1111/j.1751-5823.2010.00112.x. URL https://doi.org/10.1111/j.1751-5823.2010.00112.x.

[45] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A "Nutrition Label" for Privacy. *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009. doi: 10.1145/1572532.1572538. URL http://portal.acm.org/citation.cfm?doid=1572532.1572538.

[46] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. A Conundrum of Permissions: Installing Applications on an Android Smartphone. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7398 LNCS:68–79, 2012. ISSN 03029743.

[47] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Privacy as Part of the App Decision-Making Process. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, page 11, 2013. doi: 10.1145/2470654.2466466. URL http://dl.acm.org/citation.cfm?doid=2470654.2466466.

[48] Jennifer King. Taken out of Context: An Empirical Analysis of Westins Privacy Scale. *Symposium on Usable Privacy and Security*, pages 1–8, 2014.

[49] Jennifer King and South Hall. "How Come I'm Allowing Strangers To Go Through My Phone?" Smartphones and Privacy Expectations. *Symposium on Usable Privacy and Security*, 2012. ISSN 1556-5068. doi: 10.2139/ssrn.2493412.

[50] Jennifer King, Airi Lampinen, and Alex Smolen. Privacy: Is There an App for That? *SOUPS '11: Proceedings of the Seventh Symposium on Usable Privacy and Security*,

pages 1–20, 2011. doi: http://doi.acm.org/10.1145/2078827.2078843.

[51] Bart P. Knijnenburg. Information Disclosure Profiles for Segmentation and Recommendation. *Symposium on Usable Privacy and Security*, pages 4–7, 2014.

[52] Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. Preference-based Location Sharing: Are More Privacy Options Really Better? *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, page 2667, 2013. doi: 10.1145/2470654.2481369. URL http://dl.acm.org/citation.cfm?id=2470654.2481369.

[53] Stefan Korff and Rainer Böhme. Too Much Choice: End-User Privacy Decisions in the Context of Choice Proliferation. *SOUPS '14: Proceedings of the Tenth Symposium on Usable Privacy and Security*, pages 69–87, 2014. URL https://www.usenix.org/conference/soups2014/proceedings/presentation/korff.

[54] D Krane, L Light, and D Gravitch. Privacy On and Off the Internet: What Consumers Want. Technical report, Harris Interactive, 2002.

[55] Ponnurangam Kumaraguru and Lorrie Faith Cranor. Privacy Indexes: A Survey of Westin's Studies. Technical Report December, 2005.

[56] Bryan W Lewis. The irlba Package. pages 1–7, 2015. URL https://cran.r-project.org/web/packages/irlba/vignettes/irlba.pdf.

[57] Jialiu Lin, Norman Sadeh, Shahriyar Amini, Janne Lindqvist, Jason I. Hong, and Joy Zhang. Expectation and Purpose. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 501–510, 2012. ISSN 1617-4909. doi: 10.1145/2370216.2370290.

[58] Jialiu Lin, Michael Benisch, Norman Sadeh, Jianwei Niu, Jason Hong, Banghui Lu, and Shaohui Guo. A Comparative Study of Location-Sharing Privacy Preferences in the United States and China. *Personal and Ubiquitous Computing*, 17(4):697–711, 2013. ISSN 16174909. doi: 10.1007/s00779-012-0610-6.

[59] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I. Hong. Modeling Users Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. *Proceedings of the tenth Symposium on Usable Privacy and Security*, 1:1–14, 2014.

[60] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, page 271, Denver, CO, 2005. USENIX Association. ISBN 9781931971317. URL https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu.

[61] Bin Liu, Jialiu Lin, and Norman Sadeh. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help? *Proceedings of the 23rd International Conference on World Wide Web*, pages 201–212, 2014. doi: 10.1145/2566486.2568035. URL http://dl.acm.org/citation.cfm?id=2566486.2568035.

[62] Bin Liu, Bin Liu, Hongxia Jin, and Ramesh Govindan. Efficient Privilege De-Escalation for Ad Libraries in Mobile Apps. *MobiSys*, pages 89–103, 2015. doi: 10.1145/2742647.2742668.

[63] Minxing Liu, Haoyu Wang, Yao Guo, and Jason Hong. Identifying and Analyzing the Privacy of Apps for Kids. *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*, pages 105–110, 2016. doi: 10.1145/2873587.2873597. URL http://doi.acm.org/10.1145/2873587.2873597.

[64] Mary Madden, Lee Rainie, Andrew Perrin, Maeve Duggan, and Dana Page. Americans' Attitudes About Privacy, Security and Surveillance, 2015. ISSN 15232409. URL http://www.pewinternet.org/2015/05/20/americans-attitudes-about-privacy-security-and-surveillance/.

[65] Miguel Malheiros, Soeren Preibusch, and M. Angela Sasse. "Fairly truthful": The Impact of Perceived Effort, Fairness, Relevance, and Sensitivity on Personal Data Disclosure. In Michael Huth, N Asokan, Srdjan Čapkun, Ivan Flechais, and Lizzie Coles-Kemp, editors, *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 7904 LNCS, pages 250–266. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. ISBN 9783642389078. doi: 10.1007/978-3-642-38908-5{_}19.

[66] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15(4):336–355, 2004. ISSN 10477047. doi: 10.1287/isre.1040.0032.

[67] Jonathan Mugan, Tarun Sharma, and Norman Sadeh. Understandable Learning of Privacy Preferences Through Default Personas and Suggestions. Technical report, 2011.

[68] Mohammad Nauman, Sohail Khan, and Xinwen Zhang. Apex: Extending Android Permission Model and Enforcement with User-Defined Runtime Constraints. *ASIACCS*, pages 328–332, 2010. ISSN 9781605589367. doi: 10.1145/1755688.1755732.

[69] H Nissenbaum. Privacy as Contextual Integrity. *Wash. L. Rev.*, pages 101–139, 2004. ISSN 09700420. doi: 10.1109/SP.2006.32. URL http://heinonlinebackup.com/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79&section=16.

[70] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *The Journal of Consumer Affairs*, 41 (1):100–126, 2007. ISSN 1745-6606. doi: 10.1111/j.1083-6101.2009.01494.x.

[71] Mike Orcutt. Personal AI Privacy Watchdog Could Help You Regain Control of Your Data, 2017. URL https://www.technologyreview.com/s/607830/personal-ai-privacy-watchdog-could-help-you-regain-control-of-your-data/.

[72] Raja Parasuraman, Thomas B. Sheridan, and Christopher D. Wickens. A Model for Types and Levels of Human Interaction with Automation. *IEEE transactions on Systems, Man,*

and Cybernetics. Part A, Systems and humans : a Publication of the IEEE Systems, Man, and Cybernetics Society, 30(3):286–297, 2000. ISSN 1083-4427. doi: 10.1109/3468.844354.

[73] Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J. Lee. Reflection or Action?: How Feedback and Control Affect Location Sharing Decisions. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, pages 101–110, 2014. doi: 10.1145/2556288.2557121.

[74] Paul Pearce, Adrienne Porter Felt, Gabriel Nunez, and David Wagner. AdDroid. *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security - ASIACCS '12*, page 71, 2012. doi: 10.1145/2414456.2414498. URL http://dl.acm.org/citation.cfm?id=2414456.2414498.

[75] Sai Teja Peddinti, Allen Collins, Aaron Sedley, Nina Taft, Anna Turner, and Allison Woodruff. Perceived Frequency of Advertising Practices. In *Symposium on Usable Privacy and Security*, 2015.

[76] Soeren Preibusch. Managing Fiversity in Privacy Preferences: How to Construct a Privacy Typology. 2014. URL http://cups.cs.cmu.edu/soups/2014/workshops/privacy/s1p3.pdf.

[77] J Ross Quinlan. *C4.5: Programs for Machine Learning*. Elsevier, 2014.

[78] Amir Rahmati and Harsha V. Madhyastha. Context-Specific Access Control: Conforming Permissions With User Expectations. *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 75–80, 2015. doi: 10.1145/2808117.2808121. URL http://dl.acm.org/citation.cfm?id=2808117.2808121.

[79] Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman Sadeh. Capturing Social Networking Privacy Preferences. *Symposium on Usable Privacy and Security*, page 1, 2009. ISSN 1617-4909. doi: 10.1145/1572532.1572587. URL http://portal.acm.org/citation.cfm?doid=1572532.1572587.

[80] Matthew Rogers. LBE Privacy Guard Monitors and Controls What Permissions Your Android Apps Have, 2011. URL http://lifehacker.com/5807797/lbe-privacy-guard-monitors-and-controls-what-permissions-your-android-apps-have.

[81] Cynthia E Schairer, Cynthia Cheung, Caryn Kseniya Rubanovich, Mildred Cho, Lorrie Faith Cranor, and Cinnamon S Bloss. Disposition Toward Privacy and Information Disclosure in the Context of Emerging Health Technologies. *Journal of the American Medical Informatics Association*, 2019. ISSN 1527-974X. doi: 10.1093/jamia/ocz010. URL https://academic.oup.com/jamia/advance-article/doi/10.1093/jamia/ocz010/5426084.

[82] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A Design Space for Effective Privacy Notices. *Proceedings of the Eleventh Symposium on Usable Privacy and Security*, pages 1–17, 2015.

[83] Kristi Lane Scott. United States Department of Justice Overview of the Privacy Act of 1974, 2015. URL https://www.justice.gov/opcl/file/793026/download.

[84] Scott Thurm and Yukari Iwatani Kane. Your Apps Are Watching You, 2011. URL http://www.wsj.com/articles/SB10001424052748704694004576020083703574602.

[85] Fuming Shih, Ilaria Liccardi, and Daniel J Weitzner. Privacy Tipping Points in Smartphones Privacy Preferences. *Proc. of the 2015 ACM Conference on Human Factors in Computing Systems*, pages 807–816, 2016. doi: 10.1145/2702123.2702404.

[86] Daniel J Solove. Privacy Self-Management and the Consent Dilemma. *Harvard Law Review*, page 1880, 2013. URL http://scholarship.law.gwu.edu/faculty_publications.

[87] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. *EC '01 Third ACM Conference on Electronic Commerce*, pages 38–47, 2001. doi: 10.1145/501158.501163.

[88] Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. The Effect of Developer-Specified Explanations for Permission Requests on Smartphone User Behavior. *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing systems*, pages 91–100, 2014. doi: 10.1145/2556288.2557400. URL http://dl.acm.org/citation.cfm?id=2556288.2557400.

[89] Humphrey Taylor. Most People Are "Privacy Pragmatists" Who, While Concerned About Privacy, Will Sometimes Trade It Off for Other Benefits, 2003. URL https://theharrispoll.com/wp-content/uploads/2017/12/Harris-Interactive-Poll-Research-Most-People-Are-Privacy-Pragmatists-Who-While-Conc-2003-03.pdf.

[90] US Secretary's Advisory Committee. Records, Computers and the Rights of Citizens. 1973. URL https://aspe.hhs.gov/report/records-computers-and-rights-citizens.

[91] Jeroen van de Hoven, Martijn Blaauw, Wolter Pieters, and Martijn Warnier. Privacy and Information Technology, 2014. URL http://plato.stanford.edu/entries/it-privacy/.

[92] Na Wang, Bo Zhang, Bin Liu, and Hongxia Jin. Investigating Effects of Control and Ads Awareness on Android Users Privacy Behaviors and Perceptions. *ACM MobileHCI*, pages 373–382, 2015. doi: 10.1145/2785830.2785845. URL http://dl.acm.org/citation.cfm?id=2785845.

[93] Alan Westin. Harris-Equifax Consumer Privacy Survey 1991. 1991. URL https://www.semanticscholar.org/paper/Harris-Equifax-Consumer-Privacy-Survey-1991-Westin/2036dd51d811cf4bd999313104488e8b80cc5cb4#paper-header.

[94] Alan Westin. *Equifax-Harris Consumer Privacy Survey 1996*. Equifax, 1996. URL https://books.google.com/books?id=hzpEAQAAIAAJ.

[95] Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. Privacy Manipulation

and Acclimation in a Location Sharing Application. *UbiComp*, page 549, 2013. doi: 10.1145/2493432.2493436.

[96] Pamela Wisniewski, a K M Najmul Islam, Bart P Knijnenburg, and Sameer Patil. Give Social Network Users the Privacy They Want. *Proc. of the 2015 Computer-Supported Cooperative Work and Social Computing*, (October):1427–1441, 2015. doi: 10.1145/2675133.2675256.

[97] Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Laura Brandimarte, and Alessandro Acquisti. Would a Privacy Fundamentalist Sell Their DNA for 1000...If Nothing Bad Happened as a Result? The Westin Categories, Behavioral Intentions, and Consequences. *SOUPS '14: Proceedings of the Tenth Symposium on Usable Privacy and Security*, pages 1–18, 2014. URL https://www.usenix.org/conference/soups2014/proceedings/presentation/woodruff.

[98] Guo Xun Yuan, Chia Hua Ho, and Chih Jen Lin. Recent Advances of Large-Scale Linear Classification. *Proceedings of the IEEE*, 2012. ISSN 00189219. doi: 10.1109/JPROC.2012.2188013.

[99] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman Sadeh, Steven M Bellovin, and Joel Reidenberg. Automated Analysis of Privacy Requirements for Mobile Apps. In *NDSS 2017*, volume 3078, 2017. ISBN 1891562460.

[100] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N. Cameron Russell, and Norman Sadeh. MAPS: Scaling Privacy Compliance Analysis to a Million Apps. *Proceedings on Privacy Enhancing Technologies*, 3:66–86, 2019. URL https://usableprivacy.org/static/files/popets-2019-maps.pdf.