Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions

Bin Liu,* Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi
Shikun Zhang, Norman Sadeh,* Alessandro Acquisti, Yuvraj Agarwal
Carnegie Mellon University
Pittsburgh, PA, USA
{bliu1, manderse, fschaub, hazim, shikunz, sadeh, yuvraj.agarwal}@cs.cmu.edu
acquisti@andrew.cmu.edu

ABSTRACT

Modern smartphone platforms have millions of apps, many of which request permissions to access private data and resources, like user accounts or location. While these smartphone platforms provide varying degrees of control over these permissions, the sheer number of decisions that users are expected to manage has been shown to be unrealistically high. Prior research has shown that users are often unaware of, if not uncomfortable with, many of their permission settings. Prior work also suggests that it is theoretically possible to predict many of the privacy settings a user would want by asking the user a small number of questions. However, this approach has neither been operationalized nor evaluated with actual users before. We report on a field study (*n*=72) in which we implemented and evaluated a Personalized Privacy Assistant (PPA) with participants using their own Android devices. The results of our study are encouraging. We find that 78.7% of the recommendations made by the PPA were adopted by users. Following initial recommendations on permission settings, participants were motivated to further review and modify their settings with daily "privacy nudges." Despite showing substantial engagement with these nudges, participants only changed 5.1% of the settings previously adopted based on the PPA's recommendations. The PPA and its recommendations were perceived as useful and usable. We discuss the implications of our results for mobile permission management and the design of personalized privacy assistant solutions.

1. INTRODUCTION

Mobile app ecosystems such as Android or iOS compete in part based on the number, and the quality, of apps they offer. To attract developers and help generate more apps, these platforms have exposed a growing number of APIs. These APIs provide access to smartphone functionality (e.g., GPS, accelerometer, camera) and user data (e.g., unique identifiers, location, social media accounts), much of which is privacy-sensitive.

While the Android and iOS platforms both rely on permission-based mechanisms and allow users to control access to sensitive data and functionality, the end result is an unwieldy number of app-permission decisions that users are expected to make. Estimates indicate that users, on average, have to make over one hundred permission decisions (95 installed apps on average per user [48]; 5 permissions on average per app [37]). Prior work has shown that users are often unaware of – if not uncomfortable with – many of the permissions they have ostensibly consented to at some point (e.g., [6,8,16,17,21,24]).

To help overcome the burden associated with managing such a large number of decisions, prior research suggests that – despite the diversity of users' privacy preferences – it is theoretically possible to predict many of a user's permission settings by asking the user a small number of questions [28, 29]. These approaches suggest that, using machine learning, it may be possible to reduce user burden when it comes to configuring mobile app permission settings. However, this approach has not been fully operationalized so far.

We propose a practical solution that operationalizes privacy preference modeling in a personalized privacy assistant (PPA) by (1) developing privacy profiles for users, (2) determining which of these profiles is the best match for a given user, and (3) configuring many of the user's permissions based on the selected profile. This paper is the first to report on the implementation and field evaluation of a personalized privacy assistant (PPA) for mobile app permissions.

We propose a methodology to learn privacy profiles for permission settings and leverage these profiles in a personalized privacy assistant that actively supports users in configuring their permission settings. In a field study we collected permission settings from 84 Android users with rooted smartphones who received privacy nudges designed to motivate them to interact with their permission settings. Mobile app permission settings collected from these users were organized along three dimensions: app categories, app permissions and purposes associated with each permission (e.g., supporting an app's core functionality versus advertising). The resulting data was used to identify clusters of like-minded users and to generate recommended permission settings (or "profiles") for users in each cluster. Our results indicate that despite relying on app permission settings collected from a small number of users (*n*=84), our learned privacy profiles can accurately recommend mobile app permission settings that users are likely to adopt.

Our personalized privacy assistant uses information about the apps installed on a user's smartphone to elicit the user's privacy preferences and offer recommendations on how to configure associated

*Main contacts: Bin Liu and Norman Sadeh.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22–24, 2016, Denver, Colorado.

permission settings. We designed an interactive profile assignment dialog, in which the PPA relies on dynamically-generated decision trees to generate questions that help match users to the privacy profile that best aligns with their preferences, which is then used to provide recommendations on which permissions to deny. The PPA gives the user the option to accept multiple recommended settings at once and the ability to modify them as needed.

We show the effectiveness and usability of a profile-based PPA through a field study. The profiles built using permission settings collected from the first set of users (*n*=84) were used by our PPA, which we evaluated in a second between-subjects field study with different participants (*n*=72). This enabled us to evaluate the effectiveness and usability of the PPA on participants' own (rooted) Android smartphones. Our results show that 78.7% of the recommendations made by the PPA were accepted by participants in the treatment group, and only 5.1% of recommended permission settings were later revised by participants, despite being exposed to privacy nudges designed to motivate them to revisit their earlier decisions. Participants in the treatment group also converged faster on their settings and reported satisfaction with the recommendations and the PPA functionality.

Our results provide rich insights on the interaction design of personalized privacy assistants, permission managers, mobile privacy nudges, and their interplay. These insights are relevant for developers of mobile platforms, privacy tools, and mobile apps.

2. RELATED WORK

Our work relates to research on mobile privacy, mobile app permissions, privacy awareness, and building privacy profiles for users.

2.1 Mobile App Privacy

Prior work has shown that many mobile apps access sensitive functionality and data for purposes that are not limited to the delivery of their core functionality [5, 13, 27, 49]. Sensitive resources and data commonly accessed by mobile apps, whether on iOS or Android, include unique device identifiers (e.g., IMEI), user location, contacts list, camera, texting, and much more. Many apps share sensitive personal information with advertising networks and analytics companies, which in turn use the data to build extensive user profiles [1,34,47,49]. Research shows that users are often unaware of the extent of these practices and that many will express reservations and concern when they learn about them [18, 23, 25, 27, 45].

2.2 App Privacy Management

Functionality that enables users to manage mobile app permissions has evolved quite significantly in recent years – for both iOS and Android. While early versions of iOS only allowed users to control access to their location, the number of such permissions has increased in each new version of iOS. In iOS 9, 11 categories of permissions exist with settings enabling users to grant or deny individual permissions on an app-by-app basis, at the time the permission is requested by an app. Until recently, the user privacy controls provided by Android were fairly limited. They mainly involved displaying a list of permissions to the user when installing an app and asking the user to confirm that they consent to grant all the requested permissions. In Android 6.0, this has changed, with both Android and iOS now offering very similar control over mobile app permissions to their users. While this increase in control is a positive development, it also exposes users to a large number of privacy settings.

Prior work has shown that mobile app permission screens at install time are largely ineffective in helping users make informed privacy decisions, because most users do not pay close attention to the permissions screen and do not understand what the permissions mean or entail [16, 23]. Alternative designs that highlight privacy implications (e.g., how personal information is shared with advertisers [24] or unexpected data collection practices [27]) have been more effective in helping users avoid what they perceive as intrusive apps [9, 21, 24, 27, 35, 50]. Instead of assisting decisions about whether to install an app, our work focuses on helping users manage their privacy for apps already installed on their devices.

In Android 6.0, Google replaced install-time permission screens with just-in-time permission requests and a permission manager [7], reminiscent of iOS' permission management approach. Prior work has explored the utility and usability of such permission managers showing how users employ them to limit app access to personal information [6, 19]. Fisher et al. found that the majority of iOS users in their study prevented a third of their apps from accessing the users' location [19]. Similarly, Almuhimedi et al. found that 65% of Android users in their study utilized the permission manager to control how apps access personal information [6]. However, they also showed that the permission manager alone is not sufficient for users to reach satisfying levels of privacy protection because the permission manager does not provide enough information to assist users in making informed privacy decisions [6]. To account for such a limitation, we enrich the permission manager in our study with additional information such as the purpose and access frequency information for specific permissions.

Both iOS and Android 6.0 encourage app developers to specify a purpose in permission request dialogs in order to enable users to make informed privacy decisions. Tan et al. evaluated the prevalence of such developer-specified explanations in iOS apps (only 19% of permission requests had explanations) and observed that while users did not really understand them they were still more likely to grant requests if an explanation was provided [46]. Using experience sampling, Shih et al. find an opposite effect: participants shared more when permission requests did not contain explanations, whereas vague explanations decreased users' willingness to grant permission requests [44]. Instead of relying on developer-specified explanations, we notify users of the likely purpose of an app's permission request, based on static code analysis results from PrivacyGrade [2,27,28]. Prior work indicated that purpose explanations play an important role in making privacy decisions [6,27,44].

A number of recent studies explored approaches to help users manage their privacy for apps they already installed on their devices [6, 8, 20]. Fu et al. showed in a field study that a full-screen and interruptive privacy notification is more effective than an uninterruptive icon in the notification area in informing users when apps access their location [20]. However, users were annoyed by the full-screen notifications, especially when apps accessed location frequently [20]. Using just-in-time notifications when personal information is accessed and a summary of how frequently apps access users' information, Balebako et al. showed that users are in general unaware of data collection practices by apps and that users are surprised at how frequently apps access their personal information [8]. Both Fu et al. and Balebako et al. did not provide users with tools to exercise control over how apps access users' personal information. In contrast, we enabled our users to manage their app privacy settings through an enhanced permission manager. To explore whether interventions can motivate users to review their app privacy settings, Almuhimedi et al. designed "privacy nudges" that inform users of how frequently apps access personal information (e.g., location), and also enable users to adjust their app settings [6].

They found that nudges indeed increase awareness of apps' behaviors and motivate users to review and adjust their app permissions.

In this paper, we build on some of the ideas proposed by prior work. In particular, in addition to showing frequency of access to private data, we also show the inferred purpose of the access using the public PrivacyGrade dataset [2]. Second, while we build upon the idea of privacy nudges, we extend it to elicit user preferences on a set of privacy-related questions to build privacy profiles with machine learning. Finally, we build on prior work on using privacy profiles to reduce user burden in terms of decisions, but we extend it to use privacy nudges to help users review their settings after profile assignment to ensure that profile-based settings match users' actual preferences. Most importantly, our PPA app integrates these aspects in an end-to-end system to evaluate their effectiveness in real-world settings.

2.3 Privacy Profiles and Preference Modeling

Privacy controls, such as permission managers, enable users to configure their privacy settings. However, the growing number of configurable privacy settings makes it difficult for users to align their privacy settings with their actual preferences [6, 32] Agarwal and Hall [5] and Rashidi et al. [39] proposed crowd-powered and expert-powered systems to recommend settings to users. However, users' app privacy settings are diverse [29], rendering one-size-fits-all solutions insufficient to accurately capture users' diverse preferences.

Researchers have proposed modeling and predicting users' privacy preferences. Collaborative filtering has been proposed for location sharing preferences [53, 54]. However, the proposed approaches were only evaluated in simulations. In real-world scenarios for mobile apps, the collaborative filtering solutions would suffer from data sparsity and the cold-start problem, where the model requires sufficient user feedback before giving accurate recommendations. Ismail et al. [22] proposed a collaborative-filtering-based recommender for security configurations of mobile apps. They determined a sufficiency threshold for user input before providing recommendations. And they pre-determined diverse scenarios for users to ensure informativeness of the training input.

Privacy profiles, which are collections of related privacy and sharing rules that correspond to privacy preferences of similar-minded users [11, 15, 26, 28, 29, 40, 51, 52], can provide decision support if one can identify a privacy profile that matches a new user. In the context of online social networks, Fang and LeFevre suggested using active machine learning to design a "privacy wizard" to assist Facebook users in managing their complex privacy settings [15]. The authors evaluated the privacy wizard using real data from 25 Facebook users and showed that the privacy wizard can predict users' privacy settings with high accuracy (above 90%) and minimal effort by users (only labeling 25 friends) [15]. In the context of mobile app privacy, recent work has explored utilizing related approaches. Lin et al. [28] generated privacy profiles for app privacy settings, taking into consideration purpose information and users' self-reported willingness to potentially grant access, elicited in a scenario-based online study. However, the privacy paradox suggests that self-reported preferences may not necessarily reflect actual privacy behavior [10, 31]. In contrast, Liu et al. identified six privacy profiles based on 239K real users using only their app privacy settings [29]. However, prior work shows that permission settings alone might not reflect users' actual privacy preferences, because users may be unaware of many apps' data collection practices occurring in the background [6]. In contrast, we built privacy profiles from users' real-world permission settings collected in a field study using permission settings, purpose information as well as app categories to obtain a diverse set of profiles from a comparatively smaller dataset. We further use privacy nudges to make users aware of unexpected data practices and thus elicited privacy settings likely better aligned with users' privacy preferences.

In contrast to prior work, we evaluated the effectiveness of our privacy profiles with actual users in a field study, thereby, demonstrating the practical impact of privacy profiles on mobile privacy configuration. Few others have evaluated privacy profiles on real users' phones in the field. Wilson et al. studied privacy profiles in the context of a location-sharing system [51]. They found that privacy profiles impacted users' privacy decisions and satisfaction level. However, they evaluated their privacy profiles based on simulated location requests, whereas we evaluated our privacy profiles based on real permission requests on participants' own smartphones.

3. PPA OVERVIEW

We designed and implemented a profile-based personalized privacy assistant (PPA).[1] Specifically, the PPA uses apps on the user's smartphone to engage in a dialog and elicit a small set of preferences pertaining to whether or not the user feels comfortable granting some permissions to apps from certain categories. Using these answers, the PPA identifies a privacy profile that best matches the user's preferences and, based on this profile, recommends a number of permission settings changes to the user. The user is given the option to accept or change recommendations individually or in bulk. The specific set of questions the PPA asks a user is determined by the user's installed apps and dynamically adapts as the user answers questions.

Developing and deploying our PPA involved multiple steps. We first collected users' app privacy preferences using an enhanced permission manager on rooted Android devices to develop mobile app privacy preference profiles. We organized users into clusters of like-minded people, and developed profiles for each cluster to capture typical user preferences. Next, a field study was conducted where we deployed the PPA to newly recruited users, also with rooted Android devices. In this study, the PPA used its profiles to engage in dialogs with users and assign them to a particular cluster. The profiles were finally used to recommended specific mobile app permission settings to users. This is further detailed below.

Enhanced Android Permission Manager
For the purpose of accurately capturing users' privacy preferences from their privacy settings, we assume that users are comfortable with a restrictive permission setting they chose, if they keep the setting and do not change it back to a permissive setting. 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 [12], which we describe below.

Simplified controls. In the permission manager, we organized permission settings into six groups of privacy-related permissions: Location, Contacts, Messaging, Call Log, Camera, and Calendar. As a result, multiple permissions are represented as a single permission, reducing the overall number of permissions users have to consider. For example, `READ_CONTACTS` and `WRITE_CONTACTS` are represented as "Contacts." This grouping is partially based on results by Lin et al. [27] and Felt et al. [16]. Users can directly allow or deny

[1] Our personalized privacy assistant app is publicly available at: www.privacyassistant.org

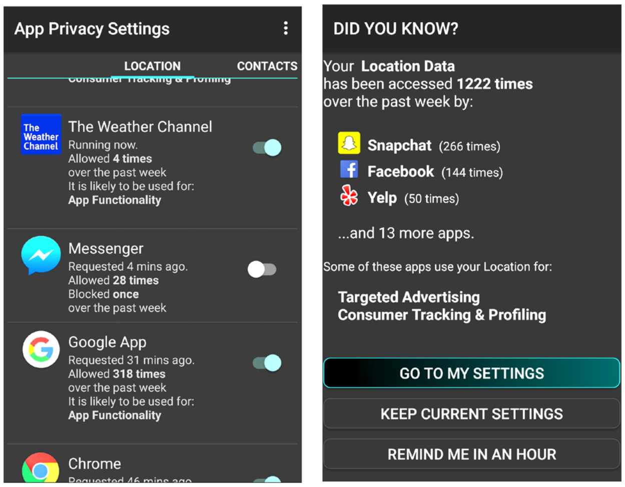

The Weather Channel
Running now.
Allowed **4 times**
over the past week
It is likely to be used for:
App Functionality

Messenger
Requested 4 mins ago.
Allowed **28 times**
Blocked **once**
over the past week

Google App
Requested 31 mins ago.
Allowed **318 times**
over the past week
It is likely to be used for:
App Functionality

Chrome

Your **Location Data**
has been accessed **1222 times**
over the past week by:

Snapchat (266 times)
Facebook (144 times)
Yelp (50 times)

...and 13 more apps.

Some of these apps use your Location for:

Targeted Advertising
Consumer Tracking & Profiling

Figure 1: Permission manager (*left*) and a daily privacy nudge (*right*), which include the access frequency and purpose information.

each permission while reviewing them in the permission manager.[2]

Enhanced Awareness. We extended the permission manager to show not only an app's most recent access requests, but also how often the app requested access over the last seven days, as shown in Figure 1. We further included purpose information from PrivacyGrade [2,28] for apps for which it was available. Using Androguard static analysis [27], PrivacyGrade identifies the likely purpose(s) of an app's permission requests by analyzing its third-party libraries (e.g., app functionality, targeted advertising, consumer tracking & profiling, or sharing with social network services).

Privacy Nudges. Nudges have been found to be effective at increasing users' privacy awareness and motivating them to review and adjust their permissions [6, 9]. We adopt a similar nudging strategy to get users to reflect on their permissions and engage with our permission manager to adjust their settings, in order to collect rich permission settings from each user. Our privacy nudge, shown in Figure 1, includes access frequency for the given permission [6], other apps that accessed the same permission, and, if known, the likely purpose of the access for that permission. From the nudge, users can open the permission manager to change their settings, keep the current settings and close the nudge, or postpone managing their privacy.

Building Profiles

After deploying our enhanced permission manager to users, we collect their real-world permission settings. For each permission setting, we collect the likely purpose of the permission request from PrivacyGrade [2], and the category of the requesting app from the Google Play store. We use app categories as features, rather than individual apps, to reduce over-fitting caused by less popular apps and limited training samples. Using this training data, we build user profiles by applying hierarchical clustering [43] on the feature vectors generated from a set of features. We describe the process of building privacy profiles from real users' privacy settings in more detail in Section 4.

Assigning users to privacy profiles

In order to assign new users to the generated privacy profiles, we ask them a small number of tailored questions about their privacy preferences. To generate these questions, we first aggregate user preferences in the training data set by (a) each permission; (b) each (permission, app category) pair; and (c) each (permission, purpose) pair. Each aggregated feature represents a potential question to ask a new user. However, we first check whether users have apps installed that fit the particular question. For example, to be asked a question about preferences for (location, advertisement), the user must have at least one app installed that accesses location for advertisement purposes. We then train a C4.5 decision tree [38] on the set of questions applicable to a particular user, and generate an ordered list of questions. Users are asked 5 questions at most to be assigned to a profile. Note that with our method the set of questions is dynamically personalized for each user, so that the questions can be contextualized using the apps each user has installed on their phones.

Generating recommendations

On the server side, we train a scalable SVM classifier (LibLinear [14]) using the permission settings we collected from the profile-building procedure mentioned above. The PPA app will pass the user's features to the classifier to generate recommendations for privacy settings learned from the training data. The features we include are the user's assigned profile, app category, permission, and purposes. Even though our model can make recommendations for each (category, permission, purpose) tuple, Android's permission model does not support granular control by purposes. Therefore, our personalized privacy assistant provides privacy recommendations to deny access based on permission and app categories, while we use purpose information to further explain our recommendations. Note that we only provide recommendations to deny access, as permissions were allowed by default once an app was installed prior to Android 6.0.

Next, we discuss our process for building privacy profiles in Section 4, followed by a discussion of the design of our personalized privacy assistant in Section 5.

4. BUILDING PRIVACY PROFILES

To obtain real users' permission settings from which to build privacy profiles, we conducted a first field study in which we deployed our enhanced permission manager to actual Android users.

4.1 Privacy Settings Dataset Collection

Since permission management requires system privileges, this study (as well as the later evaluation of our PPA) 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. In previous online surveys and studies using dialogs on simulated phone screens [28, 50], settings selected by participants were not applied to devices actually used by these participants. In contrast, our approach allows us to collect real settings stemming from user behavior, rather than aspirational responses that don't match users' behavior [31]. 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 PPA, 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.

[2] Coincidentally, Google announced similarly grouped permissions for Android 6.0 shortly after we conducted our first field study.

Our study was approved by Carnegie Mellon University's Institutional Review Board. We recruited Android phone users (>1 month use) 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. 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. 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 user name to activate it.

In the first week of the study, participants could use the permission manager to selectively deny or allow permissions. 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 12pm and 8pm. Figure 1 shows both the permission manager (left) and the nudge dialog (right). 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 was shown was randomized to counter order effects. If a particular permission had never been accessed by apps on the participant's device (access frequency would be zero), another permission would be selected to be shown in the nudge instead.

After participants completed the study, we asked them to fill an exit survey online, consisting of the 10-item IUIPC scale on privacy concerns [30] and an 8-item scale on privacy-protective behavior [36]. They were compensated with a $15 giftcard afterwards. 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, and the usability of the enhanced permission manager and privacy nudges. Those who participated in the optional interview received an additional $10 giftcard.

4.2 Dataset Analysis

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. 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 had an Associate's degree; 30 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 [30] 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). They also took more measures to protect their online privacy compared to the general population [36], as shown in Table 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 of the 6 permissions in the enhanced permission manager. We filtered the dataset to only analyze permission settings for apps available in the Google Play Store. Because Android permission requests of installed apps are set to allow by default,[3] 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.

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, 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. The independent variables include major features that could be obtained in our dataset such as user demographics and app category. App category information was retrieved from the Google Play store. The detailed logistic regression results are shown in Table 2 in Appendix A. App category and the type of permission are significant predictors for an individual's allow or deny decision, whereas demographics, privacy concerns, the app name, access frequency and purpose information were not significant.

Participants largely 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. Participants' aggregated 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 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.

Nine interviewees (out of ten) confirmed the usefulness of access frequency information; four stated it was as a reason to deny a permission, five mentioned it was useful in the nudge, and two stated

[3] All participants use Android 4.4.X or 5.X phones, where app permissions were granted by default when an app is installed. Android 6 prompts users to grant or deny permission requests, thus making this pre-processing unnecessary.

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 significant impact of access frequency on users’ decision of permission settings (see Table 2).

The logistic regression model indicates that purpose information was not a significant predictor for whether a permission is denied in our dataset. A likely reason is the sparsity of purpose information compared to app category and permission type which are always available. Our purpose information stems from PrivacyGrade’s dataset [2], which covers popular free apps on Google Play. During the study, purpose information was shown 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% contained purpose information; 31.45% of those nudges showed purposes other than required for app functionality. Participants denied less if any purpose(s) were shown (13.53% compared to 19.95%; Chi-square=10.1793, df=1, p=0.0021, effect size(odds ratio)=0.6784), which matches Tan et al.’s results [46]. However, none of the purposes had significant impact on users’ decisions (see Table 2). Participants further agreed on some specific cases. For instance, 100% allowed Contacts for Social Network Services and 95.63% allowed Camera for App Functionality. Nine interviewees mention that purpose information was useful; three as a reason to deny, seven as useful in the nudge, and three as useful in the permission manager. Three interviewees mentioned a trade off 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 additional purpose information is useful to participants and it would be desirable to provide it for more apps. However, it seems some purposes also caused confusion. P3 had problems understanding the meaning of “Consumer Tracking / Profiling.” Thus, more research is needed to reliably determine purposes of permission requests, convey this information to users, and enable users to make access decisions for specific purposes. We discuss these aspects in more detail in Section 7.2.

4.3 Generating Privacy Profiles

From the collected dataset, we obtained users’ detailed app permission settings as a collection of rows in the form of (user, app, permission, decision). We collected app category information from the Google Play store. Purpose information is based on PrivacyGrade data [2], 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.

4.3.1 Clustering Approach

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 permissions 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. To estimate similarities among participants’ feature tensors, we impute the missing values in the tensors. In order to impute without biasing any dimension, we apply weighted PARAFAC Tensor factorization [3]. We put 1-weight on all known data cells and 0-weight on unknown data cells in the tensor. Thus, we optimize the overall error of the imputed tensor in Frobenius norm using only the values known from the data. Using the users’ feature vectors reshaped from the imputed tensor, we build user profiles by applying hierarchical clustering [43] on the feature vectors. We choose hierarchical clustering since it is not sensitive to the size or density of clusters and allows non-Euclidean distances.

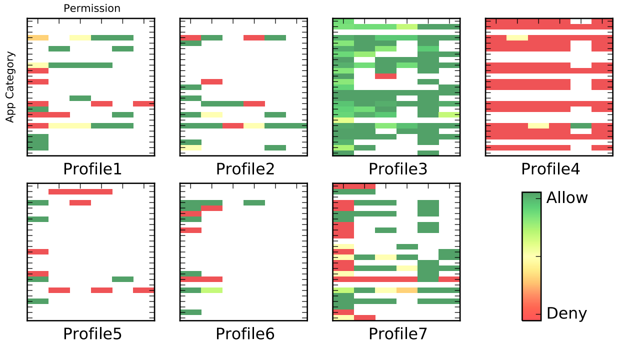

Figure 2: Privacy profiles learned from collected app privacy settings. Profile 1 is more protective on Location and Productivity apps than other profiles. Profile 2 denies phone call log permission more. Profile 3 is generally permissive. Profile 4 denies most permission requests. Profile 5 generally denies contacts, message, phone call log and calendar access, with only location and camera allowed for some apps. Profile 6 denies location and contact access of Social apps and Finance apps. Profile 7 is stricter regarding Social apps and location access in general.

4.3.2 Generating Recommendations

The profile-based recommended settings are generated by a scalable SVM Classifier (LibLinear [14]) 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.

4.3.3 Resulting privacy profiles

We applied a grid-search of the parameters for the hierarchical clustering and the SVM classifier to choose the ones that have better cross-validated F-1 scores of the accuracy of the recommended items to deny. We tried Manhattan, Euclidean, and Cosine distances in the grid search of parameters for hierarchical clustering, and tried Gamma={0,1e-3, 1e-4} and C={1e-4, 1e-3, ..., 1e3} for the linear-kernel SVM. With 5-fold cross-validation on the dataset described in Section 4.2, we found the optimized mode for the dataset (hierarchical clustering: K=7, complete linkage, cosine distance, Silhouette Coefficient=0.2079; classifier: Gamma=1e-3, C= 1e3, hinge loss) with a cross-validated F-1 score of 90.02%. In contrast, if we train a global model for all users without splitting them into profiles, the best F-1 score would be 74.24%, much lower than the profile-based optimized model.

Figure 2 shows the permission preferences in each profile aggregated by app categories. It provides an overview of the diversity in privacy preferences among the different profiles. Profile 3 contains 67 of the 84 participants (79.8%), who are generally permissive. Profile 4 contains 2 participants (2.4%), who denied most permission requests. Note that the majority of participants were grouped in the most permissive profile (profile 3) despite our privacy-conscious and tech-savvy participant population. The remaining profiles (15 participants, 17.8%) express variations in privacy preferences depending on app category and permission of ac-

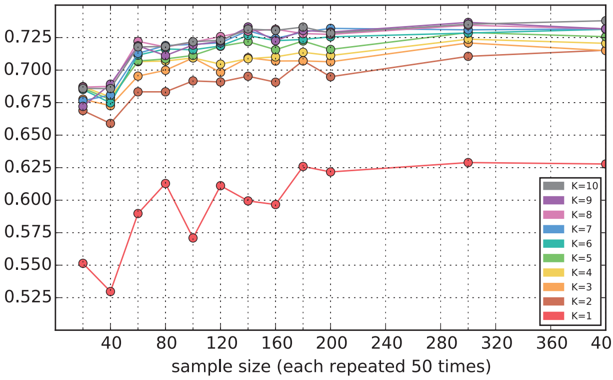

Figure 3: Down-sampling simulation on Lin et.al's dataset [28] (F-1 score). With 5 profiles or more training on data from just 80 users provides reasonable F-1 score (> 70%). When training on 400 users, the accuracy improves, but only marginally.

cess. Profile 1 (3 participants) is more protective on Location and on apps in the category of Productivity comparing to other profiles. Profile 2 (4) denies phone call log permission more. Profile 5 (1) generally denies contacts, message, phone call log and calendar permission access to all apps, with only location and camera allowed for some. Profile 6 (3) denies location and contact access of Social apps and Finance apps. Profile 7 (4) is restrictive for Social apps and location access in general.

Lin et al. [28] identified similar profiles. Their "unconcerned" profile corresponds to our profile 3, their "conservative" profile to profile 4, and their "fence-sitter" and "advanced users" profiles align with our more specialized profiles (profiles 1, 2, 5, 6, 7).

4.3.4 Downsampling comparison

Given the relatively small number of 84 participants in our dataset, a potential concern is whether our profiles are expressive enough to cover privacy preferences of a larger user population, and whether we can provide useful recommendations. To explore the utility of our profiles, we applied our approach for building profiles to Lin et al.'s considerably larger dataset [28]. This dataset has 21,657 records in total, consisting of 725 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 randonmly-selected users to create smaller datasets, ranging from 20 to 400 users in size, which is more than half of the entire dataset. Figure 3 shows F-1 scores for 1–10 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 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.

5. PROVIDING RECOMMENDATIONS

Our PPA app elicits a user's privacy preferences with an interactive dialog to provide the user with personalized recommendations. Thus, the PPA's recommendation process consists of two main components: (a) First, the PPA shows a series of dynamically-generated questions to elicit the user's app privacy preferences and assign the user to a privacy profile. (b) Then, the PPA provides profile-based recommendations according to the user's privacy profile and installed apps. The user can review and adjust recommended settings before applying them.

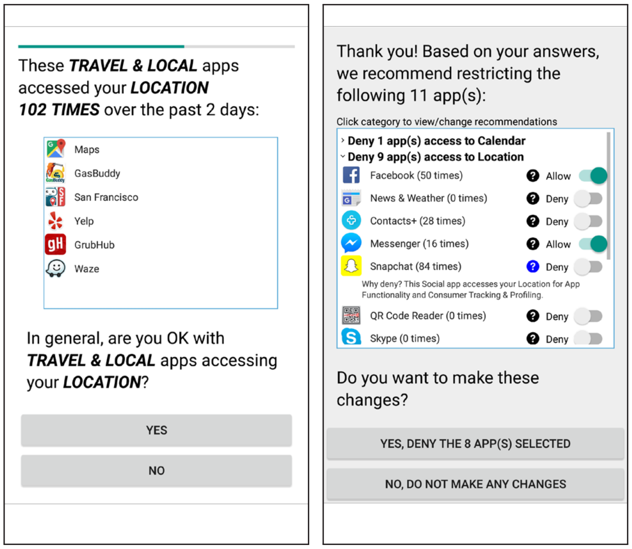

These ***TRAVEL & LOCAL*** apps accessed your ***LOCATION*** ***102 TIMES*** over the past 2 days:

- Maps
- GasBuddy
- San Francisco
- Yelp
- GrubHub
- Waze

In general, are you OK with ***TRAVEL & LOCAL*** apps accessing your ***LOCATION***?

YES

NO

Thank you! Based on your answers, we recommend restricting the following 11 app(s):

Click category to view/change recommendations

› **Deny 1 app(s) access to Calendar**

˅ **Deny 9 app(s) access to Location**

Facebook (50 times)	Allow	●
News & Weather (0 times)	Deny	○
Contacts+ (28 times)	Deny	○
Messenger (16 times)	Allow	●
Snapchat (84 times)	Deny	○

Why deny? This Social app accesses your Location for App Functionality and Consumer Tracking & Profiling.

QR Code Reader (0 times)	Deny	○
Skype (0 times)	Deny	○

Do you want to make these changes?

YES, DENY THE 8 APP(S) SELECTED

NO, DO NOT MAKE ANY CHANGES

Figure 4: Profile assignment dialog: After answering up to 5 questions (*left*) users may receive personalized recommendations (*right*). Users can review and customize the recommended deny settings.

5.1 Interactive Profile Assignment

The profile-assignment questions elicit a user's preferences for (1) individual permissions, (2) permission and app category pairs, and (3) permission and purpose pairs. Each question has a Yes/No response. For a new user, the PPA dynamically generates a decision tree that uses input from a question to determine the next question to ask and eventually assign the user to one of our privacy profiles. Users are asked 5 questions at most to be assigned to a profile. The decision tree is generated based on profile assignments and aggregated preferences from the dataset used to build the privacy profiles, as well as the user's installed apps. Considering installed apps allows us to contextualize the decision tree by excluding questions for which the user has no apps installed. For example, if the user has no Game app installed, the PPA would not ask if the user would generally allow Game apps to access location.

To contextualize the questions in the profile assignment dialog, installed apps that fit the particular question are listed in the dialog with their access frequency for the respective permission, inspired by Almuhimedi et al.'s privacy nudges [6]. Figure 4 shows an example of an assignment dialog question. In this example, installed apps from the Travel & Local category have accessed the Location permission 102 times over the past 2 days. A progress bar at the top shows how many questions have been completed.

5.2 Profile-based Recommendations

After a user has responded to the questions, the PPA assigns a privacy profile to the user, which is used to determine which recommendations to show. For each permission requested by apps on the user's phone, the PPA applies the classifier trained with the profiles (see Section 4.3.2) to generate an allow/deny decision for the user. The PPA will then display a list of recommended restrictive permission changes to the user.

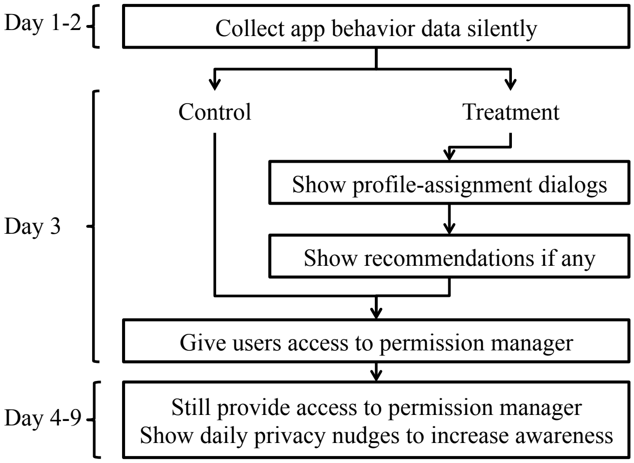
Day 1-2
Collect app behavior data silently
Control
Treatment
Day 3
Show profile-assignment dialogs
Show recommendations if any
Give users access to permission manager
Day 4-9
Still provide access to permission manager
Show daily privacy nudges to increase awareness

Figure 5: Overview of the study protocol for the two conditions.

Recommendations are grouped by permission (e.g., Calendar, Location); these groups can be expanded to view individual apps, as shown in Figure 4. For each app, clicking the question mark reveals an explanation for this specific recommendation, referencing the user's responses to the profile assignment questions. For instance, in Figure 4 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 input of the user's profile and the apps' characteristics, such as its category and the purpose of permission requests.

6. FIELD STUDY: EVALUATING THE PPA

We conducted another field study with a second group of Android users with rooted devices to evaluate the effectiveness of our privacy profiles in the context of our PPA. In this study, 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 profile-based support. Participants in both conditions had access to our enhanced permission manager and received privacy nudges.

6.1 Study Procedure

We wanted to evaluate the effectiveness of the profile-based PPA with participants from the same population the privacy profiles were based on. Hence, we followed the same recruitment approach as in the data collection study. We extended the screening survey to exclude individuals with prior experience using other Android permission or privacy managers. We also excluded any participants from our first study. After qualifying for the study, the newly-recruited participants received a user id and instructions for installing the study client.

Our study protocol is summarized in Figure 5. During day 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 at 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, as described in Section 5.1. Users were assigned to a profile and personalized recommendations were generated, as described in Section 5.2. 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 match their preferences or if they 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. 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.

6.2 Results

We received valid screening survey responses from 138 participants. We excluded 4 participants who had participated in the first study and 3 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 first study suggested that many participants may have permissive privacy attitudes, in which case they may be assigned to profile 3 (most permis-

Table 1: Privacy protective measures of our study populations compared to the general population. Questions and general population results are based on a Pew survey [36].

Population	Pew Survey	Data Coll. 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%
Given 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%

sive) and thus would not receive restrictive recommendations and, hence, would not interact with the recommendation screen (shown on the right in Figure 4). Thus, we increased the number of treatment participants to account for these considerations.

6.2.1 Demographics

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 U.S.), 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 [30]): 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). The participants' measures to protect their online privacy compared to the general online population [36] are shown in Table 1.

6.2.2 Effectiveness of recommendations

In the treatment group, the number of received recommendations depended on the privacy profile participants were assigned to and their installed apps. 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 got assigned to Profile 3, the most permissive profile. Another participant was assigned to Profile 2 but did not have any of the apps installed that were denied in the assigned privacy profile.

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 (they were from profile 1 (4 of them), 2(3), 3(6) and 7(2)), 9 accepted some (they were from profile 1(2), 2(2), 5(3) and 7(2)), and 3 accepted none (all from profile 3; they were shown only one recommendation). Figure 6 shows the number of accepted and rejected recommendations for each of these participants.

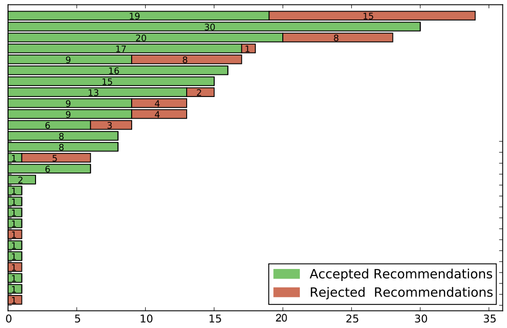

Figure 6: The numbers of recommendations accepted or rejected by participants receiving them. Overall, users accept 78.7% of all recommendations.

The 15 participants that accepted all recommendations primarily stated that they did so because the recommendations matched their preferences (11) or that they trusted the PPA (8). Note that participants could provide multiple reasons. The 3 participants that accepted no recommendations stated that it would have restricted app features (3) or broken app functionality (1), or that the recommendations did not reflect their preferences (2). The 9 participants who accepted some recommendations also stated restricted (6) or broken (4) app functionality as a reason for non-acceptance; 4 stated the recommendations did not reflect their preference, while only 1 responded that they did not like that the PPA wanted to change so many settings automatically.

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 the recommendations were effective (high precision).

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 7. Among the 383 permission settings changes made by the treatment group, the participants made 316 (82.51%) of them during day 3, which is the day when they received profile-based recommendations and the first day when they had access to the permission manager. In contrast, the control group only made 68.42% (104 of 152) of their permission settings on day 3. The difference of the treatment and the control condition has 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 have 43 respective 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

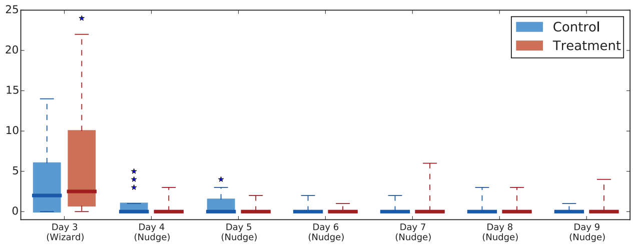

Figure 7: Number of permission changes in the control and treatment groups on the different days of the study. On day 3, the treatment group got recommendations; and both groups were given access to the permission manager.

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, 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. 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, and then compare 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 allowed Camera access more. 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 for which purposes permissions are granted forced them to be more permissive than desired, i.e., they lack the capabilities to regulate privacy as desired.

Participants are comfortable with provided recommendations. 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. 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 8. 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.

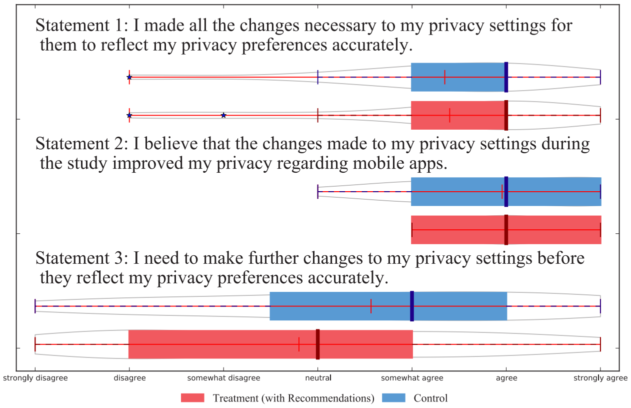

Figure 8: Participants’ responses about their privacy settings in the exit questionnaire. Participants who received recommendations felt slightly less of a need to make further changes to their settings.

6.2.3 *Usability of the personalized privacy assistant*

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).

Nudge timing and delivery is important. When asked about what they liked the least, participants from both conditions identified timing of the nudges as an issue (18 treatment, 13 control). Asked how we could improve the PPA, participants from both groups suggested to turn 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 have 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 trade-offs 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.

7. DISCUSSION

Our results suggest that personalized privacy assistants can indeed help users better manage their mobile app permission settings. They provide evidence based on deployment with actual users that profile-based recommendations can help users configure their mobile app permissions. Below, we first discuss limitations of our work, followed by insights gained about the development and interaction design of personalized privacy assistants.

7.1 Limitations

Because manipulating people's mobile app permission settings requires root access, the target population available for recruitment for this study was limited. As a result, the sample populations in both filed studies skew young, male, tech-savvy, and privacy-conscious. Accordingly, one might expect the privacy settings and permission profiles obtained for this population to be more conservative (namely, more restrictive) than those of the general population. But one cannot be entirely sure: rooted users are also more technically sophisticated and possibly more daring. In fact, a relatively large number of our participants selected rather permissive privacy settings. 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 main objective was to evaluate (1) a practical approach for collecting permission data and learning profiles, and (2) 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. 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, Samsung, or Apple, could implement our approach (or provide APIs for researchers and developers) and support functionality similar to the one evaluated in this study.

In contrast to prior work, we learned privacy profiles from a relatively small dataset, which could be viewed as a limitation. We overcame this potential limitation by collecting rich, real-world permission data and aggregating obtained permission settings along three dimensions, namely app category, permissions, and purpose information. Our second field study validates the effectiveness of the learned profiles and recommendations. Three-quarters (78.7%) of the provided recommendations were accepted, and only a small number of recommendations to restrict permissions were later re-allowed (5.1%) – primarily because the restrictive permissions impaired some app functionality, rather than participants having privacy preferences that differed from those in the assigned profiles. Participants further reported high comfort with their privacy settings at the end of the study.

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 fairly 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. This data was used to learn privacy profiles and provide participants in the second study with privacy recommendations to support initial configuration. The low number of subsequent permissions changes (see Figure 7) furthers support the notion that PPA users had converged on stable settings by the end of the study. 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.

7.2 Privacy Profiles and Recommendations

Our results show the feasibility of learning privacy profiles from a relatively small number of users. These profiles are effective at supporting users in configuring their permission settings and helping them make privacy decisions. In the second field study, which evaluated the profile-based PPA, participants reviewed and accepted 78.7% of our recommendations. Additionally, very few recommended restrictive permission settings were changed back by participants (5.1%). However, some participants restricted additional permissions based on information shown in the privacy nudges and the permission manager. This suggests that our classifier could possibly be tuned to provide more aggressive recommendations. It is also likely that having access to a larger corpus of permission settings would enable us to build profiles with higher predictive power. Finally, the ability to directly adjust recommended settings and the option to make additional changes in the permission manager was perceived as useful by most participants, as it helped them reflect on their privacy settings and bootstrap the configuration.

Our recommendations could further be improved with enhanced filtering techniques to exclude core system apps and services, as well as apps that crash when restricted. App crashes were sometimes reported as a reason for re-allowing permissions. 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, since app developers will adapt and add exception handling for denied permissions.

A general issue that emerged was a conflict between restrictive privacy preferences and permissions required by an app to properly function. This happens when apps require permissions for multiple purposes (e.g., both to support their core functionality and to support advertising). 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, as it would have broken essential features of the application. This 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 [44, 46], once access is granted, apps can currently 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 like Google Play Services, it is effectively impossible for users to make meaningful decisions about granting or denying Google Play access to a permission 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 controls that better align with users' privacy decision making. While these finer-grained models could increase user burden, our research suggests that they may in fact lend themselves to the learning of more powerful predictive models, which in turn could actually help reduce user burden by providing a larger number of more accurate recommendations.

For future personalized privacy assistants, we envision to assist users with privacy monitoring, configuration, and decision support beyond initial permission configuration. Settings recommendations could be provided when installing new apps or as part of just-in-time permission requests. Ultimately, privacy assistants should further adapt to users by learning their privacy preferences over time, for instance by engaging with them in a continuous, yet unobtrusive, dialog. Micro-interactions initiated at opportune times and tailored to the user's context [41, 42] could help increase the usability of privacy nudges by better integrating them into a user's interaction flow. This also requires enhancing machine learning techniques to appropriately account for the uncertainty, contextual nature, and malleability of privacy preferences [4].

7.3 Designing Personalized Privacy Assistants

Our two field studies provided extensive insights on how users interact with different mobile privacy tools: our enhanced permission manager, privacy nudge interventions, privacy profile assignment dialogs, 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, and should therefore be taken into consideration when designing personalized privacy assistants and the associated user experience.

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. 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. In addition to using access frequency of the installed apps, we plan to explore the utility of creating statistical models of how often specific apps access certain resources in order to be able to provide permission recommendations without a training phase. This information could in addition be added to an app's app store information, enabling users to use frequency in decision making even before installing an app.

Privacy recommendations introduce a degree of automation to privacy configuration. Automation can potentially impact technology acceptance [33]. 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.

Our results show that the enhanced privacy manager – including both information on permission access frequency and purpose – helped participants monitor app behavior and manage their privacy settings effectively. 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. Furthermore, many participants mentioned the nudge's timing and modality as an issue. However, the use of modal dialogs was a conscious choice to force interaction with the nudge messages in our study. In the public release version of our PPA, we implemented nudges as standard Android notifications to make them less obtrusive.

While our results and insights pertain primarily to mobile interaction, we expect that personalized privacy assistant approaches can also be applied to support privacy decision making in other domains where privacy configuration or awareness is an issue. For instance, in the context of websites, where privacy policies are often difficult to understand, or the Internet of Things (IoT), where secondary channels will have to be utilized for privacy management, because most IoT devices have small or no screens [41].

8. CONCLUSION

In this paper, we demonstrated how users can benefit from a personalized privacy assistant that provides them with recommendations for privacy configuration. Our personalized privacy assistant is based on privacy profiles learned from real-world permission settings. Our proposed approach is practical and can learn representative privacy profiles even from a relatively small number of users (*n*=84). We evaluated the effectiveness of the privacy profiles by conducting a field study (*n*=72), in which we deployed our personalized privacy assistant on participants' own smartphones (rooted Android devices). Our results show that 78.7% of recommendations were accepted by users and that only 5.1% of settings were changed back during the study. Overall, the assistant led to more restrictive permission changes without sacrificing users' comfort with these settings.

9. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under grants CNS-1012763 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 grand from the CMU-Yahoo! InMind project, as well as by the Carlsberg Foundation. The views and conclusions contained herein are those of the authors 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!

The authors would like to thank the anonymous reviewers and our shepherd Marian Harbach for their constructive feedback.

10. REFERENCES

[1] Android Flashlight App Developer Settles FTC Charges It Deceived Consumers. https://goo.gl/Zf18jI, 2013. Accessed: 2016-02-01.

[2] PrivacyGrade: Grading The Privacy Of Smartphone Apps. http://privacygrade.org, 2015. Accessed: 2016-02-01.

[3] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup. Scalable tensor factorizations with missing data. In *SDM*, pages 701–712. SIAM, 2010.

[4] A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, Jan. 2015.

[5] Y. Agarwal and M. Hall. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proc. MobiSys*, 2013.

[6] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. Cranor, and Y. Agarwal. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proc. CHI*. ACM, 2015.

[7] arstechnica. Android M Dev Preview delivers permission controls, fingerprint API, and more. http://goo.gl/NdmOx1, 2015. Accessed:2016-02-01.

[8] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen. Little brothers watching you: Raising awareness of data leaks on smartphones. In *Proc. SOUPS*, 2013.

[9] E. K. Choe, J. Jung, B. Lee, and K. Fisher. Nudging people away from privacy-invasive mobile apps through visual framing. In *Proc. INTERACT*, 2013.

[10] K. Connelly, A. Khalil, and Y. Liu. Do i do what i say?: Observed versus stated privacy preferences. In *Proc. INTERACT 2007*, pages 620–623. Springer, 2007.

[11] J. Cranshaw, J. Mugan, and N. Sadeh. User-controllable learning of location privacy policies with gaussian mixture models. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.

[12] EFF. Awesome Privacy Tools in Android 4.3+. https://www.eff.org/deeplinks/2013/11/awesome-privacy-features-android-43, 2013. Accessed: 2015-2-17.

[13] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information flow tracking system for real-time privacy monitoring on smartphones. *Comm. ACM*, 2010.

[14] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874, 2008.

[15] L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proc. WWW '10*. ACM, 2010.

[16] A. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android Permissions: User Attention, Comprehension, and Behavior. In *Proc. SOUPS '12*, 2012.

[17] A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proc. CCS '11*, pages 627–638. ACM, 2011.

[18] A. P. Felt, S. Egelman, and D. Wagner. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proc. SPSM*, 2012.

[19] D. Fisher, L. Dorner, and D. Wagner. Short paper: location privacy: user behavior in the field. In *Proc. SPSM '12*, pages 51–56. ACM, 2012.

[20] H. Fu, Y. Yang, N. Shingte, J. Lindqvist, and M. Gruteser. A field study of run-time location access disclosures on android smartphones. In *Proc. USEC*, 2014.

[21] M. Harbach, M. Hettig, S. Weber, and M. Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proc. CHI*, 2014.

[22] Q. Ismail, T. Ahmed, A. Kapadia, and M. K. Reiter. Crowdsourced exploration of security configurations. In *Proc. CHI '15*, pages 467–476. ACM, 2015.

[23] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall. A conundrum of permissions: installing applications on an android smartphone. In *Proc. FC '12*. Springer, 2012.

[24] P. G. Kelley, L. F. Cranor, and N. Sadeh. Privacy as part of the app decision-making process. In *Proc. CHI*, pages 3393–3402. ACM, 2013.

[25] J. King. How come i'm allowing strangers to go through my phone? smartphones and privacy expectations. In *Proc. SOUPS*, 2013.

[26] B. P. Knijnenburg. Information disclosure profiles for segmentation and recommendation. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*, 2014.

[27] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proc. UbiComp*, 2012.

[28] J. Lin, B. Liu, N. Sadeh, and J. I. Hong. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Proc. SOUPS*, 2014.

[29] B. Liu, J. Lin, and N. Sadeh. Reconciling mobile app privacy and usability on smartphones: could user privacy profiles help? In *Proc. WWW '14*. ACM, 2014.

[30] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (iuipc): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.

[31] P. A. Norberg, D. R. Horne, and D. A. Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.

[32] L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *Proc. CHI '03*, pages 129–136. ACM, 2003.

[33] R. Parasuraman, T. Sheridan, and C. D. Wickens. A model

for types and levels of human interaction with automation. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 30(3):286–297, May 2000.

[34] Path official blog. We are sorry. `http://blog.path.com/post/17274932484/we-are-sorry`, 2012. Accessed:2016-02-01.

[35] A. Paturi, P. G. Kelley, and S. Mazumdar. Introducing privacy threats from ad libraries to android users through privacy granules. In *Proc. USEC '15*. Internet Society, 2015.

[36] Pew Research Center. Internet project/GFK privacy panel. `http://www.pewinternet.org/files/2015/05/Privacy-and-Security-Attitudes-5.19.15_Topline_FINAL.pdf`, 2014. Accessed:2016-02-01.

[37] Pew Research Center. An Analysis of Android App Permissions. `http://www.pewinternet.org/2015/11/10/an-analysis-of-android-app-permissions/`, 2015. Accessed:2016-02-01.

[38] J. R. Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.

[39] B. Rashidi, C. Fung, and T. Vu. Dude, ask the experts!: Android resource access permission recommendation with recdroid. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 296–304, May 2015.

[40] R. Ravichandran, M. Benisch, P. G. Kelley, and N. M. Sadeh. Capturing social networking privacy preferences. In *Proc. PET '09*, pages 1–18. Springer, 2009.

[41] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A design space for effective privacy notices. In *Proc. SOUPS '15*, pages 1–17, Ottawa, July 2015. USENIX Association.

[42] F. Schaub, B. Konings, and M. Weber. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *Pervasive Computing, IEEE*, 14(1):34–43, Jan 2015.

[43] Scikit-Learn. Scikit-learn manual. `http://scikit-learn.org/stable/modules/generated/sklearn.cluster.AgglomerativeClustering.html`. Accessed:2016-02-01.

[44] F. Shih, I. Liccardi, and D. J. Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proc. CHI*. ACM, 2015.

[45] I. Shklovski, S. D. Mainwaring, H. H. Skúladóttir, and H. Borgthorsson. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proc. CHI*, 2014.

[46] J. Tan, K. Nguyen, M. Theodorides, H. Negrón-Arroyo, C. Thompson, S. Egelman, and D. Wagner. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proc. CHI*. ACM, 2014.

[47] The Guardian. Uber faces FTC complaint over plan to track customers' locations and contacts.

[48] The Next Web. Android users have an average of 95 apps installed on their phones, according to Yahoo Aviate data. `http://thenextweb.com/apps/2014/08/26/android-users-average-95-apps-installed-phones-according-yahoo-aviate-data/#gref`, 2014. Accessed:2016-02-01.

[49] S. Thurm and Y. I. Kane. Your apps are watching you. `http://www.wsj.com/articles/SB10001424052748704368004576027751867039730`, 2010. Accessed: 2016-02-01.

[50] N. Wang, B. Zhang, B. Liu, and H. Jin. Investigating effects of control and ads awareness on android users' privacy behaviors and perceptions. In *Proc. MobileHCI '15*. ACM, 2015.

[51] S. Wilson, J. Cranshaw, N. Sadeh, A. Acquisti, L. F. Cranor, J. Springfield, S. Y. Jeong, and A. Balasubramanian. Privacy manipulation and acclimation in a location sharing application. In *Proc. UbiComp '13*, pages 549–558. ACM, 2013.

[52] P. Wisniewski, B. P. Knijnenburg, and H. Richter Lipford. Profiling facebook users' privacy behaviors. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*, 2014.

[53] J. Xie, B. P. Knijnenburg, and H. Jin. Location sharing privacy preference: Analysis and personalized recommendation. In *Proc. IUI '14*, pages 189–198. ACM, 2014.

[54] Y. Zhao, J. Ye, and T. Henderson. Privacy-aware location privacy preference recommendations. In *Proc. Mobiquitous '14*, 2014.

APPENDIX

A. LOGISTIC REGRESSION RESULTS

Results of the random effect logistic regression are shown in Table 2.

Table 2: Random effect logistic regression on users' allow/deny decisions grouped by users (Likelihood ratio test of $\rho = 0$: $\bar{\chi}^2 = 338.10$, $P >= \bar{\chi}^2 : 0.000$).

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**
App Category	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		