To Deny, or Not to Deny: A Personalized Privacy Assistant for Mobile App Permissions

[Draft]

Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Norman Sadeh, Yuvraj Agarwal, Alessandro Acquisti
Carnegie Mellon University
Pittsburgh, PA, USA
Contact email: sadeh@cs.cmu.edu

ABSTRACT

Many smartphone users are uncomfortable with the permissions requested by their mobile apps. The sheer number of permissions can be so overwhelming that many users are unable to adequately manage their permission settings. We present a methodology for building personalized privacy assistants to recommend permission settings to users. We conducted two field studies with Android users: the first (n=84), to collect privacy preferences and build a recommendation system, the second (n=51), to evaluate the effectiveness of the recommendations. Results show that 73.7% of recommendations are accepted. Following interactions with the assistant, participants were motivated to further review and modify their settings with daily privacy nudges. Despite showing significant engagement and modifying permissions not covered in the recommendations, participants only modified 5.6% of the recommendations they had accepted. We discuss implications of our results for the design of existing permission managers and future privacy assistants.

Author Keywords

Privacy; mobile computing; nudging; privacy assistants; personalization; privacy preference profiles; Android.

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces; K.4.1. Computers and Society: Public Policy Issues—*Privacy*.

INTRODUCTION

Advanced mobile platforms such as iOS and Android, and the millions of third party apps written for them, have driven the rapid adoption of mobile devices such as smartphones and tablets. However, this rapid growth has also led to significant privacy concerns among users given the variety of privacy-sensitive resources – e.g., location, identity, contacts – apps can request access to. Given these concerns, both iOS and Android have introduced mechanisms to inform users about data accessed by apps and give them some degree of control of such access. The (in)effectiveness of these mobile privacy management and permission tools has been focus of much research (e.g., [12, 13, 16, 27, 28, 33, 36, 55, 57]), which has resulted in practical improvements. For instance, in both iOS and Android, users can now selectively restrict apps' permissions [15], and use just-in-time permission dialog [12,47] that are shown when an app requests access to a resource. Just-in-time dialog are typically complemented by a privacy or permission manager that enables users to review and adjust their settings. While these permission managers enable privacy control at a finer granularity, they also significantly increase the number of decisions users have to make. Thus, while privacy managers increase the level of control and are perceived as useful [13], in their current incarnation they are limited in their utility and effectiveness, and are becoming increasingly unmanagable. Small interventions, such as privacy nudges [8, 17], have been shown to help users reflect on their privacy settings [16] and can be effective at getting users to switch their attention from their primary task to privacy management [13]. However, mobile privacy nudges have also not seen widespread implementation so far. And, they do not reduce user burden.

Contributions

In this paper, we propose a methodology to learn privacy profiles for permission settings and use these profiles in the form of *Personalized Privacy Assistants* that actively support users in configuring their permission settings. The personalized privacy assistants leverage the apps a user has installed on his or her mobile phone to elicit their privacy preferences and offer recommendations on how to configure associated permission settings, including options to automatically configure multiple permission settings at once. To learn these profiles, we first conducted a field study in which we collected permission settings from Android users who over a period of a week had received daily nudges designed to increase their awareness of the data collected by their apps and to motivate them to revise their settings and align them with their privacy preferences. Having build profiles based on this data, we conducted a second field study with new participants to validate the effectiveness of our personalized privacy assistant in a between-subjects study. The work reported herein

FTC PrivacyCon, January 14, 2016, Washington, D.C.

makes the following contributions to research on mobile privacy and privacy preference modeling:

- We describe a practical approach for generating privacy profiles from app permission settings elicited from mobile users. We build profiles based on users' app permission preferences aggregated along three dimensions: app categories, permissions and purposes of permission requests. Using hierarchical clustering with real-world permission settings collected in a field study (n=84), we identified a diverse set of privacy profiles, which significantly boosted prediction accuracy of users' permission settings.

- To apply the privacy profiles, we developed a personalized privacy assistant for Android. We designed interactive profile assignment dialog that use dynamically-generated decision trees to match users to the privacy profile that best aligns with their preferences. Based on the matched profile, our assistant provides users with recommendations on how to adjust their app permission settings.

- We validated the effectiveness of our personalized privacy assistant in a second field study (n=51). Our results show 73.7% accuracy of user acceptance of provided recommendations, as well as more restrictive privacy settings and higher comfort with these settings compared to a control group.

- With our two field studies, we gained extensive 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. For instance, we show that enhancing app permission managers with apps' permission access frequency and purposes, improves the permission manager's utility by providing cues for privacy decision making. Enhanced permission managers and periodic nudge messages help users monitor their apps' behavior and engage in app privacy management.

RELATED WORK

Our work relates to research on mobile privacy, mobile permissions, privacy awareness, and privacy preference profiles.

Mobile Privacy and App Permissions

Prior work has shown that mobile apps access users' personal information for purposes users may not be comfortable with [2, 3, 5, 12, 23, 39, 59]. In light of such perceived privacy violations, users show strong concerns over apps' privacy practices [29, 37, 56], as well as a desire to control what types of personal information apps can access [1, 4, 13, 30].

Some research has focused on helping users make better privacy decisions when installing new apps. For instance, install-time permission screens have been shown to be ineffective [27, 35], and alternative notice designs have been proposed to help users make more informed privacy decisions when installing apps [19, 32, 33, 36, 39, 48]. In contrast, we focus on assisting users in managing their privacy in relation to apps already installed on their phone.

Previous research developed and enhanced permission managers for app privacy management [12, 18, 34, 43]. In Android 6.0, Google is replacing install-time permission screens with just-in-time permission requests and a permission manager [15], 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 [12, 13, 30], but also that permission managers alone are not sufficient for users to reach satisfying levels of privacy protection [13].

We increase the effectiveness of permission managers by enriching them with purpose and access frequency information for specific permissions. Both iOs and Android 6.0 encourage app developers to specify a purpose in permission request dialog 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 [57]. Using experience sampling, Shih et al. find the opposite [55]. Participants shared more when permission requests did not contain explanations; vague explanations decreased users' willingness to grant permission requests. Instead of relying on developer-specified explanations, we notify users about the likely purpose of an app's permission request, based on static code analysis results from PrivacyGrade [6, 39].

Privacy Nudging and Awareness

Nudges are "soft-paternalistic" behavioral interventions that aim to support user's decisions by accounting for decision making hurdles without restricting choices [8, 10, 58]. A prominent decision making hurdle in the context of privacy is asymmetric information: a gap between users' and service providers' knowledge of data practices, potential consequences, and available privacy protections [9, 10, 17]. Wang et al. proposed a Facebook privacy nudge that helps users consider the audience and content of their posts to avoid later regrets [60]. Recently, Facebook introduced a similar nudge to prevent accidentally posting publicly [24].

For mobile apps, nudging has been used to help users avoid installing intrusive apps [33, 36, 39] and, in a few cases, to support app privacy management [13, 16, 31]. Almuhimedi et al. designed privacy nudges that inform users of how frequently apps access personal information (e.g., location or contacts), and enable them to adjust their app privacy settings [13]. They find that nudges increase awareness of apps' intrusive behaviors and motivate users to review and adjust their app permissions. We utilize similar privacy nudges in both our field studies to elicit privacy preferences and settings from users. Whereas Almuhimedi et al. only showed access frequency in their nudges, we enhanced, both, privacy nudges and our permission manager, with access frequency and purpose information for specific permissions. Prior work indicated that purpose explanations play an important role in making privacy decisions [13, 39, 55].

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 [13, 45]. To help reduce user burden, researchers have proposed using privacy profiles: clusters of related privacy and sharing rules that correspond to privacy preferences of similar-minded users [21, 26, 38, 40, 41, 51, 61, 62]. Lin et al. [40] generated privacy profiles for app privacy settings, taking into consideration purpose information and users' self-reported respective willingness to potentially grant access, elicited in a scenario-based online study. However, the privacy paradox suggests that self-reported preferences may not reflect actual privacy behavior [20, 44]. In contrast, Liu et al. identified six privacy profiles based on 239K real users using only their app privacy settings [41]. 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 [13]. 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. As far as we are aware, few others have evaluated privacy profiles in the field. Wilson et al. studied privacy profiles in the context of a location-sharing system [61]. 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.

FIELD STUDY: PRIVACY SETTINGS DATA COLLECTION

Our personalized privacy assistant uses privacy profiles to recommend permission settings to users. Part of the novelty of this paper is that we created profiles based on behavioral data from actual users collected on their own devices. For this purpose, we modified and extended the Android permission manager App Ops [22] and designed an enhanced privacy nudge compared to Almuhimedi et al. [13]. Privacy nudges have been shown to increase privacy awareness and motivate users to review and adjust privacy settings. By using nudges in a two-week field study, we were able to elicit privacy behavior and capture privacy settings that are well-aligned with users' actual privacy preferences.

We enhanced the permission manager and nudges with information about purposes of specific permission requests. Purpose explanations have been shown to be a relevant factor in privacy decision making [13, 39, 55], and developer-specified explanations have been integrated into Android 6.0 and iOS' just-in-time permission requests, but we are the first to integrate purpose information into an actual permission manager, as well as nudge interventions. Further, we added access frequency information to the permission manager, because it has been shown to be a useful decision cue in nudges [13].

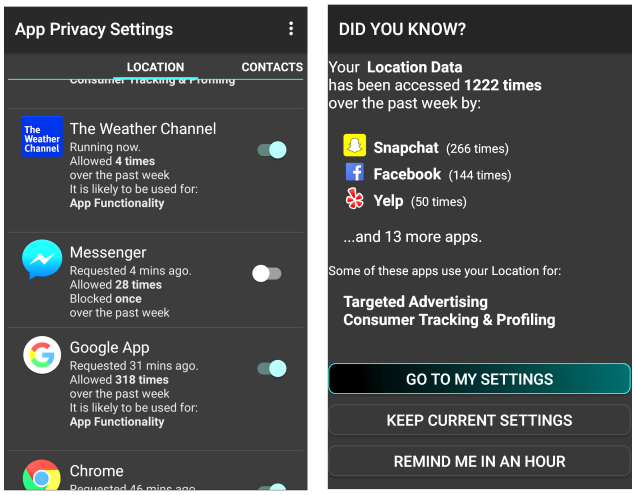

Figure 1: Permission manager (*left*) and a privacy nudge (*right*) which include the access frequency and purpose information.

Enhanced Permission Manager and Privacy Nudges

Our permission manager, shown in Figure 1, is based on Google's App Ops [22]. We heavily modified App Ops to simplify control options and enhance privacy awareness.

Simplified controls. In the permission manager, we organized permission settings into 6 groups of privacy-related permissions: Location, Contacts, Messaging, Call Log, Camera, and Calendar. As a result, multiple permissions are represented as a single permission. For example, `READ_CONTACTS` and `WRITE_CONTACTS` is represented as "Contacts." This grouping is partially based on results by Lin et al. [39] and Felt et al. [27]. Note, that Google announced similarly grouped permissions for Android 6.0 shortly after we launched our field study. In addition, we eliminated an extra interaction step by enabling users to allow/deny permission requests directly on the permission manager's main screen.

Improved Awareness. We extended the permission manager to show not only an app's most recent access request, but also how often the app requested access over the last 7 days. We further included purpose information from PrivacyGrade [6, 40]. Using static analysis [39], PrivacyGrade identifies likely purposes of an app's permission request, such as app functionality, targeted advertising, consumer tracking & profiling, and social network services.

Privacy Nudges. Almuhimedi et al. showed that nudges can effectively increase users' privacy awareness and motivate them to review and adjust their permissions [13]. We adopted a similar nudging strategy to get users to engage with the permission manager and 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 [13], examples of apps that accessed the permission, and why some of these apps likely accessed the permission. Users can open the permission manager to change their settings, keep the current settings and close the nudge, or postpone privacy management.

Study Procedure

The main goal of this study, for which we received IRB approval, was to collect permission settings from real Android users on their own smartphones. In addition, we were also interested in evaluating how our enhanced permission manager and privacy nudge designs affected participants' privacy decision making.

We first asked the participants to complete an initial screening survey to check qualifications and collect demographics information. The 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, they could use the permission manager to selectively deny or allow permissions. Our app also collected frequencies of permission requests for installed apps, which was shown in the permission manager. In the second week, the participants received a privacy nudge once a day, between 12pm and 8pm. The nudge contained information about one of six permissions: Location, Contacts, Messaging, Call log, Camera, or Calendar. The selection 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 instead. In Figure 1 right, an example of a nudge showing Location access can be seen.

After completing the study, we asked participants to complete an exit survey online, consisting of the 10-item IUIPC scale on privacy concerns [42] and a 13-item scale on privacy-protective behavior [49]. They received a $15 giftcard as compensation. We further invited them to participate in an optional interview, in which we explored (1) why participants restricted or allowed different permissions, (2) their comfort concerning their permission settings, and (3) the usability of the enhanced permission manager and privacy nudges. Interview participants received an additional $10 giftcard.

Recruitment and Demographics

We recruited participants who were 18 years or older, Android phone users (>1 month), who used a rooted Android phone (4.4.X or 5.X) with data plan. Considering that our target population was limited to users of rooted Android phones by technical requirements, 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.

In total, 114 participants passed the screening survey and installed the study app on their phones; 84 completed the study. 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 participants to skew male, young, and tech-savvy, which was the case. The majority was male (78) and 6 female. They were 18–54 years old (median 23). Among them, 8 had a graduation degree, 22 a bachelors, 5 have associates, 30 attended some college, and 19 have a high school degree or lower. The most commonly reported occupations were student (35), computer engineers or IT professionals (8), service (5), and unemployed (5). Participants exhibited relatively high privacy concern [42], scoring high for awareness (mean M=6.45, SD=.65), control (M=6.02, SD=.89), and collection (M=5.71, SD=1.17). They also took more measures to protect their online privacy compared to the general population [49]. For instance, 64.71% participants reported to disable cookies, compared to 44.16% of the general population [49].

Results

In total, we obtained 4,197 permission settings from 84 participants, reflecting their allow and deny settings at the end of the study. We filter the dataset, to only analyze permission settings for apps available in the Google Play Store. Android permission requests are allowed by default. Thus, to extract settings that reflect some level of user consent or awareness, 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 the requests. After filtering, our dataset consisted of 3,559 permission settings for 729 distinct apps.

Concerning participants' reactions to nudges, We found that participants were significantly more likely to follow to nudges mentioning "Targeted Advertising" (23.91% of occasions) compared to nudges without it (17.77%, McNemar χ^2=313.48, df=1, $p < 0.0001$). We also found that users tend to review their settings more if the nudge is about Location access (25%), compared to Messages (23.75%), Call Log (18.75%), Camera (15.19%), Calendar (14.29%), and Contacts (12.20%).

Denying app permissions

Of the 3,559 permission settings, 2,888 (81.15%) were allowed and 671 (18.85%) were denied by participants. Call Log requests were denied the most (41.33%), while Camera access was allowed the most (95.07%). Participants largely agreed on permission settings for certain app categories. For example, "Books & Reference" apps were always denied access to Contacts and Call Log, while "Photography" apps were always allowed access to Camera. For 51.14% of the (permission, category) pairs for which we have permission settings from at least three participants, the agreement among participants is at least 80%. For the remaining pairs, participants' settings are much more diverse. We calculated the standard deviation for each app category, the average deviation is 0.388.

The interviews provide insights on participants' reasons for denying apps' permissions. Nine interviewees (out of 10) confirmed the usefulness of access frequency information; 4 stated it as a reason to deny, 5 mentioned it was useful in the nudge, and 2 stated it was useful in the permission manager. For example, P1 stated: "Didn't notice that the app had actually accessed the location that many times. It is pretty crazy."

In our 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 nudges, 60.4% contained purpose information, in 31.45% of nudges shown purposes are not only app functionality. Participants generally tended to deny less (13.53% compared to 19.95%) if purposes were shown, which matches Tan et al.'s results [57]. However, participants' decisions vary by purpose: they denied 25% of permissions for Targeted Advertising, 17.65% for Consumer Tracking / Profiling, 12.68% for App Functionality, and 10% for Social Network Services. Participants agreed on some specific cases. For instance, 100% allowed Contacts for Social Network Services, 95.63% allowed Camera for App Functionality, and 50% denied Contacts for Targeted Advertising. Nine interviewees mentioned purpose information to be useful; 3 as reason to deny, 7 as useful in the nudge, and 3 as useful in the permission manager. This suggests that the additional purpose information is useful to participants. It seems some purposes caused confusion. P3 had problems understanding the meaning of "Consumer Tracking / Profiling."

In addition to denying permissions based on frequency and purpose, 8 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 cannot uninstall. Battery life was also mentioned by 4 interviewees as a reason to deny a permission, especially location. These findings align with Almuhimedi et al.'s results [13].

Not denying app permissions

Of participants' explicit permission changes, 7.58% are re-allows of permissions that they denied before. 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 2 interviewees, was that denying it broke or may 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. This is the same rationale as denying permissions because the app does not need the functionality. 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."

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." However, participants' choices were usually permissive in these cases.

DESIGNING A PERSONALIZED PRIVACY ASSISTANT

We designed and implemented a profile-based personalized privacy assistant (PPA) that consists of an interactive profile assignment dialog to (1) capture a user's preferences of app privacy settings and (2) provides personalized recommendations of app permission settings to users. In contrast to traditional recommender systems [11], which typically only provide recommendations once the system receives significant feedback about a users' preferences, we employ a profile-based approach [40, 41] to provide personalized permission recommendations with minimal interactive user input.

Our approach consists of multiple steps. First, we use the permission data obtained in the first field study to identify multiple clusters, or "privacy profiles," of like-minded users with sufficiently similar privacy preferences. Based on these profiles, we then generate a set of interactive questions shown to new users in order to assign them to a privacy profile. Finally, based on the user's installed apps and their privacy profile, we provide personalized privacy recommendations. We describe each step in further detail below.

Building profiles. First, we collect training data on real-world permission settings as described in the previous section. This training dataset includes users' settings for different requested permissions, the likely purpose of these requests, and the category of the requesting app. We use app categories as features, rather than individual apps, to reduce over-fitting. We quantify each users' preference as a 3-dimensional tensor of aggregated preferences (category, permission, purpose). The value of each tensor cell is in the range from -1 (100% deny) to 1 (100% allow). To estimate similarities among participants' feature tensors, we impute missing values. In order to impute without biasing any dimension, we apply weighted PARAFAC Tensor factorization [7], so that we can optimize the error of the imputed tensor in Frobenius norm only using the known values from the data. Using this training data, we build user profiles by applying hierarchical clustering [54] on the transformed arrays reshaped from these tensors. We choose hierarchical clustering since it is not sensitive to the size or density of clusters and allows non-Euclidean distances.

Assigning new users to profiles. In order to assign new users to the generated profiles, we ask them a few 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 dimension represents a potential question to ask a new user. However, we first check whether users have apps installed that cover 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 [50] on the set of questions applicable to a particular user, and generate an ordered list of questions using the decision tree. Users are asked 5 questions at most to be assigned to a profile. We chose C4.5 because it enables us to ask questions with optimal conditional information gain, and it allows training on

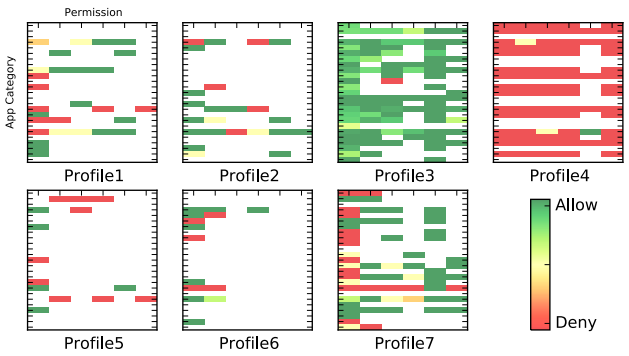

Figure 2: Privacy profiles learned based on data from the first field study. For example, Profile4 is most restrictive and Profile3 is most permissive.

sparse datasets. Note that with our method the set of questions is dynamically personalized for each user based on installed apps.

Generating recommendations. Next we train a scalable SVM classifier (LibLinear [25]) to generate recommendations for privacy settings 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 provide explanations for our recommendations.

Generated Privacy Preference Profiles

We apply a grid-search with 5-fold cross validation on the training data to choose the model parameters for hierarchical clustering and our prediction classifier. We tried Manhattan, Euclidean and Cosine distances in the grid search of parameters for hierarchical clustering. The optimized model (hierarchical clustering: K=7, complete linkage, cosine distance; Classifier: C=1e3, hinge loss) has a cross-validated training F-1 score of 90.02%, which is a substantial improvement compared to 74.24% without profiles.

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. The remaining profiles (15 participants, 17.8%) are not as polarized and express variations in privacy preferences depending on the category of the app and the purpose of access. Lin et al. [40] identified similar profiles, and labeled them similar characteristics: "unconcerned" (profile 3), "conservative" (profile 4), as well as "fence-sitter" and "advanced users" (profiles 1, 2, 5, 6, 7).

Given the relatively small number of 84 users in our dataset, a potential concern is whether our profiles are expressive enough to cover privacy preferences of a larger user population. To explore the potential benefits of larger datasets, we apply our approach for building profiles to Lin et al.'s considerably larger dataset [40], which the authors kindly provided. This dataset has 21,657 records in total, consisting of 725 MTurkers' self-reported preferences of 540 apps accessing permissions for specific purposes. We down-sample their dataset by including only smaller number of randomly selected users, ranging from 20 to 400, before applying our method. 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, the difference of prediction F-1 score becomes already less significant compared to larger sample sizes. Obviously, with training data from more users our recommendation accuracy is likely going to increase, but this experiment suggests that learning profiles from 84 field study users is still acceptable.

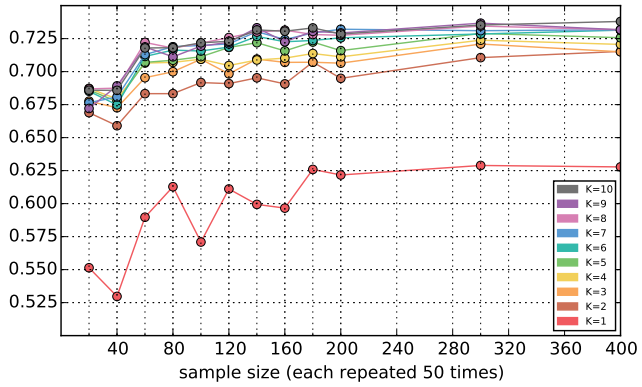

Figure 3: Down-sampling simulation on Lin et.al's dataset [40] (F-1 score). With 5 profiles or more training on data from just 80 users provides reasonable accuracy (> 70%). When training on 400 users, the accuracy improves but only marginally.

Interactive Profile Assignment Dialog

In order to evaluate our privacy profiles we have created an Android app for the PPA, with two primary UI components: (a) question screens to assign a user to a privacy profile and (b) a screen to provide recommendations to users based on their profile.

Question Screens to Assign Profiles: Our decision trees generate a series of personalized questions for each user to assign them to a profile. These questions may pertain to a permission only, permission/app-category pairs, or permission/purpose pairs. Each question has a Yes/No response. To contextualize the questions, apps that fit the particular question are list in the dialog with their access frequency for the respective permission. A representative example is shown in Figure 4. As mentioned earlier, we ask a maximum of five questions and a progress bar at the top shows how many questions have been completed.

Recommendation Screen: Once a user has been assigned to a profile, our PPA recommends restrictive permission changes for some installed apps based on the respective profile. Recommendations are grouped by permission (e.g. calendar, location); these groups can be expanded to view individual apps, see Figure 4. A "?" next to an app name can be clicked to reveal an explanation for the this specific recommendation. Based on interview feedback from the first study, we enabled users to make decisions directly on the recommen-

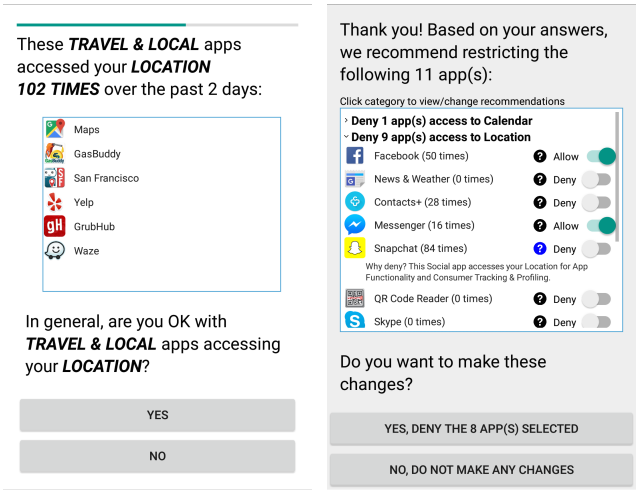

Figure 4: Profile assignment dialog: after answering up to 5 questions (*on the left*) users may receive personalized recommendations (*on the right*).

dation screen. With toggle buttons user can selectively "allow" specific recommendations, should they desire. The user can either accept all shown recommendations, accept some of them by making changes with the toggle buttons, or reject all recommendations.

FIELD STUDY: PERSONALIZED PRIVACY ASSISTANT

We conducted a second field study with android users to evaluate the effectiveness of our privacy profiles and the PPA. We also collected empirical data on how users interact with our PPA app, as well as how they modify their app permission settings. In contrast to the first field study, which primarily served the collection of training data, we conducted the second field study as a between-subjects experiment with two conditions: (a) the treatment condition in which participants interact 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 permission manager and received privacy nudges.

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 privacy managers. After qualifying for the study, the newly-recruited participants received a user id and instructions for installing the study client.

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. On the third day, the PPA initiated a dialog with participants. In the control condition, the app showed an introduction screen explaining that users could now change their settings, followed by opening the permission manager. In the treatment condition, the app also showed an introduction screen, and then initiated the profile assignment dialog, in which participants were asked up to five questions on privacy preferences, see Figure 4. Users were assigned to a profile and personalized recommendations were generated. If based on the assigned profile, no installed apps were recommended to be denied, the PPA would recommend to keep the current settings. If recommendations could be made, the recommendation screen was shown. The user could then review the list of permission recommendations 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 provide recommendations on a 7-point Likert scale, followed by a question on why they accepted all, some, or none of the recommendations. After that, the PPA also opened the permission manager to retain consistency with 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 study. During this phase, we used probabilistic experience sampling (ESM) with single-question dialogs (probability 2/3) 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 agreed to participate in exit interviews. 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, utility of the additional nudges. After completing the survey, participants were issued a $15 gift certificate as per our IRB approved protocol.

Results

Ninety-nine participants passed the initial screening survey. We excluded 4 participants who had participated in the first study and 3 participants who had prior experience with another app privacy manager. Overall, 51 participants completed the study (29 treatment, 22 control). The sample population was quite similar to the first study, with no significant differences in demographic variables, privacy concerns, or privacy-protective behavior.

Effectiveness of privacy profile based recommendations

The 29 participants in the treatment group interacted with the profile assignment dialog. They accepted 123 out of 167 provided deny recommendations (73.7%). They further denied 85 additional permissions on the same day, not covered by the recommendations, resulting in a total of 208 denied permissions. In contrast, the control group denied 101 permissions on the day the permission manager was exposed to them.

For participants in the treatment group, the number of received recommendations depended on their privacy profile and their installed apps. Thirteen of 29 participants were not shown any recommendations, either because they answered "YES" to most of the profile assignment questions or did not have any of the apps installed that were denied in their assigned privacy profile. The 16 participants who were shown recommendations, reported high comfort with the provided recommendations (Mdn.=6, M=6.21, SD=.98). This is the first indication that the profiles were helpful to users. Of the

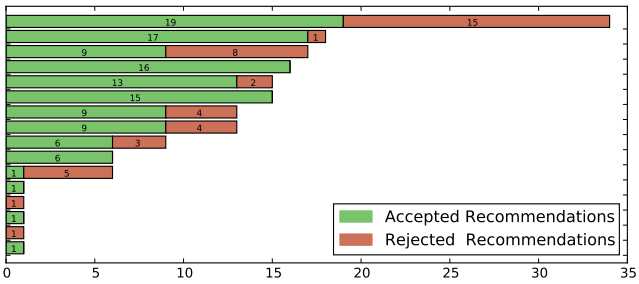

Figure 5: The numbers of recommendations accepted or rejected by participants receiving them. Overall, users accept 73.7% of all recommendations.

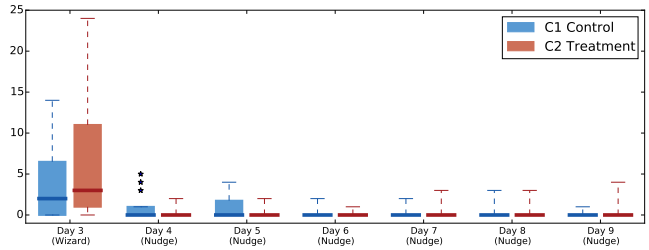

Figure 6: Number of permission changes made by the participants in the control and treatment group. On Day 3, the treatment group gets recommendations; and both groups are given access to the permission manager.

16 participants, 7 accepted all recommendations, 7 accepted some, and 2 accepted none. Figure 5 shows the number of accepted and rejected recommendations for each of these participants. A notable observation is that the 2 participants that did not accept any of the recommendation, were only shown a single recommendation.

The 7 participants that accepted all recommendations stated that they trusted the PPA (5) and that the recommendations matched their preferences (5). Note that participants could provide multiple reasons. The 2 participants that accepted no recommendations stated that it would have restricted app features (2) and 1 stated that the recommendations did not reflect their preferences. From the 7 participants that accepted some recommendations also stated restricted (5) or broken (4) app functionality as a reason for non-acceptance; 3 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.

To assess the effectiveness of the recommendations, we further analyzed the privacy behavior in the nudging phase:

Additional Denies. In the nudging phase, the treatment group denied 36 additional permissions (M=1.24, SD=1.84) and the control 39 (M=1.77, SD=2.49). We have 28 respective ESM responses from the treatment group, and 23 from control. Participants gave the following reasons for denying: "I don't use the apps features that require this permission" (treatment: 10, control: 6), "I don't want this app to use this permission" (13, 18), "The app doesn't need permission to function" (9, 11), and "Don't know" (2, 0).

Re-Allows. The number of permissions that were changed back to allow was low in both conditions (treatment: 18, M=.62, SD=1.37; control: 8, M=.36, SD=.73). This indicates that the privacy choices made on the day of recommendations, tended to be correct, and hence the recommendations were effective (high precision). Participants gave the following reasons for re-allowing : "I want to use a feature of the app that requires this permission" (ESM treatment: 2, control: 1), "I am OK with this app using this permission" (4, 1), and "The app didn't work as expected when access was restricted" (2, 1), and "Don't know" (0, 1).

The average numbers of permissions changed by participants per day of the study are shown in Figure 6. As can be seen from the figure, participants in the treatment denied more permissions on the day that they were exposed to the profile assignment dialog than the control. By the end of the study, participants in the treatment group had changed more permission settings on average (M=8.62, SD=8.43) compared to the control (M=6.72, SD=7.98). However, we found no significant difference (MannWhitney U, n.s.). Given that 13 of the 29 participants in the treatment did not see recommendations and, hence, had an experience similar to the control, we further analyzed these two subgroups: T_0 consisting of treatment participants that did receive a recommendation and T_R consisting of those participants that received at least one recommendation. A Kruskal-Wallis test showed signficant differences between T_0, T_R and the control (H(2)=9.71, p=.009). Post-hoc tests (Mann-Whitney U) showed that participants in T_R had significantly more changes than T_0 (U=38.5, p=.002) and the control (U=102, p=.015). These results can be explained by the fact that T_R consists of more privacy-protective participants, who we would expect to deny more than the more permissive group of T_0. Comfort levels with privacy settings at the end of the study, were however not statistically significant between T_R, T_0 and control (K-W, n.s.).

Usability of personalized privacy assistant

To evaluate the PPA's usability, we analyzed exit survey responses reported on a 7-point Likert scale or as open responses. When asked to state what they liked the most about the PPA, participants from both conditions reported the ability to monitor apps (13 treatment, 11 control), the app's general usability (12 treatment, 11 control). That the PPA was helpful in monitoring apps was also confirmed by treatment group participants when asked about the additional nudges (16). In terms of usability, control participants focused on the permission manager's layout and organization. Treatment participants additionally talked about the profile assignment dialog and how it supported privacy configuration. When asked about what they liked the least, participants from both conditions identified timing of the nudge as an issue (10 treatment, 13 control). One treatment participant (P5) mentioned that he wanted to have more recommendations, which supports the idea that the recommendations were useful. Asked how we could improve the PPA, participants from both groups suggested to turn the nudge into an Android notification (6 each). Treatment participants would have liked more configuration options (7), mainly to influence timing of nudges. Note that for study purposes we displayed the nudge as a modal dialog on purpose to elicit explicit interaction with the nudge.

Recommendations were useful. Out of the 29 treatment participants, 16 were shown recommendations, of whom 15 completed the exit survey. Participants found the recommendations somewhat useful (Mdn.=5, *M*=5, *SD*=1.62). The recommendations provided useful configuration support (10) and decision support (2). 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 found recommendation somewhat useful, but would prefer to set permissions manually. 3 participants found recommendations somewhat useless, stating that some recommendations would have impaired app functionality. Overall, this indicates that recommendations were useful for most participants, but better filtering of apps might improve usability.

Question dialogs were usable. Question dialogs were shown to all treatment participants. We asked them to rate on a 7-point Likert how easy or difficult the three question types were to answer. All three types were reported to be easy (permission only: Mdn=6, *M*=6.14, *SD*=.96; permission/purpose: Mdn.=6, *M*=5.96, *SD*=1.02; permission/category: Mdn=6, *M*=6, *SD*=.86). Participants also reported that the app list (*M*=6.04, *SD*=.92) and access frequency (Mdn.=6, *M*=5.43, *SD*=1.67) were useful. The app list helped create awareness of how apps used permissions (17) and to identify apps with undesired permissions (12). Access frequency also helped to improving awareness (23) and was mentioned as an important decision factor (5).

DISCUSSION

Our results show the effectiveness of our personalized privacy assistant and that our approach of learning privacy profiles from real-world permission data results in recommendations that are well accepted by different users. We first discuss limitations of our work, followed by a discussion of insights gained on the technical and interaction design aspects of personalized privacy assistants.

Limitations

In contrast to prior work, we learned privacy profiles from a relatively small dataset. 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 (73.7%) of the recommendations made by our profile-based assistants were accepted, with the vast majority of these recommendations (>94%) remaining unchanged during the following six days while participants were nudged to reconsider their decisions. Participants also reported high comfort with their privacy settings at the end of the study.

A potential limitation is the short time period of the studies. Participants may not have fully converged to stable privacy settings. We showed participants daily privacy nudges, similar to Almuhimedi et al. [13], which increased their privacy awareness and were effective at getting them to review and adjust their permission settings. We are confident that this approach enabled us to elicit permission settings for a large number of apps and permissions in a relatively short time. This data was used to learn privacy profiles and provide participants in the second study with privacy recommendations to support initial configuration. In future work, we plan to explore longitudinal interaction with personalized privacy assistants to also support continuous privacy decision making processes [14, 45, 52].

Due to the technical requirement of root access to participant's phones in order to configure app permissions, our target population for recruitment was limited. As a result, our sample populations in both studies skew towards young, male, tech-savvy, and privacy-conscious. Thus, our profiles are likely not directly applicable to other populations, but provide interesting insights into the privacy preferences of rooted phone users. We further show the validity of our approach for obtaining real-world data to learn profiles and their effective integration into a personalized privacy assistant. This approach can be applied to other target populations. Developers of apps, privacy tools, and even mobile platform providers, such as Google or Apple, could integrate our approach to learn profiles from their users' settings and provide them with personalized privacy decision support.

Insights: Privacy Profiles and Recommendations

Our results show the feasibility of learning privacy profiles from a comparatively small number of users, and that these profiles are effective at supporting users in configuring their permission settings and helping them make privacy decisions. In the second field study, participants reviewed and accepted 73.7% of our recommendations. And very few participants later-on re-allowed permission settings that were previously denied based on recommendations. However, participants often made additional denies based on information in the privacy nudges and the permission manager. This suggests that our classifier could be tuned further to provide more aggressive recommendations. At the same time, the ability to directly edit recommendations 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 downsampling experiment showed that obtaining additional permission settings data from a larger user population could further improve the effectiveness of our approach, as it would likely allow for the generation of even more comprehensive profiles requiring less additional configurations.

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 most apps will likely continue to work properly with denied permissions in the future, as is already the case on iOS, which would alleviate this issue. A larger issue was tradeoffs between restrictive privacy preferences and app functionality. 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, as it would have broken essential features of the application. Thus, current permission models need to be extended to take the purpose of permission requests into account rather than denying permissions for the whole app. While iOS and Android 6.0 support developer-specified purposes in permission requests [55, 57], once access is granted, apps can use it for any purpose. The current permission model also fails for system services, such as Google Play Services, that provide resource access (e.g., location) to multiple apps. User cannot deny them, because it is unclear which apps would be affected, yet they have no option to control which app can use these system services. An essential challenge in mobile computing 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 processes.

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. We are currently investigating how micro-interactions initiated at opportune times and tailored to the user's context [52, 53] could provide similar utility to the privacy nudges used in our studies, while better integrating them into users' interaction flow. This also requires enhancing machine learning techniques to appropriately account for the uncertainty, contextual nature, and malleability of privacy preferences [9].

Insights: 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 dialog, and profile-based recommendations. Our results show that all these tools play important, yet different, roles in supporting users with privacy configuration and decision making, and should therefore be taken into consideration when designing personalized privacy assistants and their user experience. Next, we discuss insights for each of these components.

Profile assignment is an integral part of our personalized privacy assistant, because we use a small number of privacy preference questions to provide them with privacy recommendations personalized to their installed apps. We found that participants felt confident is 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, access frequency also helped most users. However, our results indicate that app lists were most helpful in contextualizing profile assignment questions. In a real world deployment of the personalized privacy assistant, an initial training phase needed to collect access frequency statistics would likely be an obstacle to adoption. One way to counter this issue, that we plan to explore, is to create statistical models of how often specific apps access certain resources in order to provide permission recommendations from day one. This information could in addition be added to app store information about apps, enabling users to use frequency in decision making even before installing an app.

Adding privacy recommendations introduced a level of automation to privacy configuration. Automation can impact technology acceptance [46]. Our results indicate that we have achieved a good balance as participants reviewed and edited suggested permission while reporting a high level of comfort. In future work, we plan to further investigate the impact of different levels of automation on the acceptance of personalized privacy assistants. The level of automation could also become part of the personalization as well.

Our results show that enhancing the privacy manager with information on permission access frequency and purpose information led to participants make better, more informed decisions on whether or not to deny a specific permission. An improvement motivated by participants' responses, would be to include more information about how privacy and app functionality would be affected by changing the permission. It would also be interesting to extend our permission manager by adding more privacy options. This could be the possibility to limited certain app-based features (e.g., giving a banking app location access to show nearby branches but not record user location) or purpose-based restrictions, e.g., granting SnapChat access to contacts for showing contact names instead of aliases, but restricting access to contacts for user tracking and profiling, as mentioned by a participant in study 1. However, these additional features would need to be supported by the underlying permission system.

Our studies showed that adding the purpose information to a privacy nudge was useful, and nudges were found to be useful in general. Participants liked the utility of frequency and purpose information to help them monitor what apps were doing. However, many participants mentioned that the nudge's timing and modality was an issue. This was, however, a conscious choice as we wanted to ensure that participants saw, and interacted with our privacy nudges. In a public release version, this should be changed to a notification as suggested by multiple participants in both field studies.

While the results obtained focused on mobile interaction, we are confident that many of them can also be applied to support privacy decision making in other domains where privacy configuration or awareness is an issue. This is also true in the context of websites, where privacy policies are often difficult to understand, or the Internet of Things, where secondary channels will usually have to be utilized for configuration due to devices with very small or no screens.

CONCLUSION

In this paper, we demonstrate 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 which we learned from real-world permission settings where users were nudged

to align their settings with their privacy preferences. Our approach is practical and can learn representative privacy profiles even from a relatively small amount of users (n=84). We evaluated the effectiveness of the privacy profiles by conducting a second field study (n=51), in which we deployed our personalized privacy assistant on real users' own devices. Our results show that 73.7% of recommendations were accepted by users and that only 5.6% of settings were changed back during the study. Overall, the assistant led to more restrictive permission changes without sacrificing users' comfort with these settings.

ACKNOWLEDGMENTS

This material is based in part on research sponsored under NSF SBE Grant 1513957 and in part on research sponsored 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. 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!

REFERENCES

1. 2011. Futuresight: User perspectives on mobile privacy. (2011). http://www.gsma.com/publicpolicy/wp-content/uploads/2012/03/futuresightuserperspectivesonuserprivacy.pdf
2. 2012. Path official blog: We are sorry. (2012). http://blog.path.com/post/17274932484/we-are-sorry
3. 2013. Android Flashlight App Developer Settles FTC Charges It Deceived Consumers. (2013). https://goo.gl/Zf18jI
4. 2014. Pew Research Center: Public Perceptions of Privacy and Security in the Post-Snowden Era. (2014). http://www.pewinternet.org/2014/11/12/public-privacy-perceptions
5. 2015. An Update On Privacy At Uber. (2015). http://newsroom.uber.com/2015/05/an-update-on-privacy-at-uber/
6. 2015. PrivacyGrade: Grading The Privacy Of Smartphone Apps. (2015). http://privacygrade.org/home
7. Evrim Acar, Daniel M Dunlavy, Tamara G Kolda, and Morten Mørup. 2010. Scalable Tensor Factorizations with Missing Data.. In *SDM*. SIAM, 701–712.
8. A. Acquisti. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security Privacy* 7, 6 (2009), 82–85. DOI: http://dx.doi.org/10.1109/MSP.2009.163
9. Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (Jan. 2015), 509–514. DOI: http://dx.doi.org/10.1126/science.aaa1465
10. Alessandro Acquisti and Jens Grossklags. 2007. *What can behavioral economics teach us about privacy?* Taylor & Franics.
11. Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *Knowledge and Data Engineering, IEEE Transactions on* 17, 6 (2005), 734–749.
12. Yuvraj Agarwal and Malcolm Hall. 2013. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proc. MobiSys*.
13. Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. 2015. Your Location has been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proc. CHI*. ACM.
14. Irwin Altman. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. Brooks/Cole Publishing company, Monterey, California.
15. arstechnica. 2015. Android M Dev Preview delivers permission controls, fingerprint API, and more. (2015). http://goo.gl/NdmOx1
16. Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. 2013. Little brothers watching you: Raising awareness of data leaks on smartphones. In *Proc. SOUPS*.
17. Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2011. Nudging users towards privacy on mobile devices. In *Proc. CHI-PINC*.
18. Alastair R Beresford, Andrew Rice, Nicholas Skehin, and Ripduman Sohan. 2011. MockDroid: trading privacy for application functionality on smartphones. In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*. ACM, 49–54.
19. Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. 2013. Nudging people away from privacy-invasive mobile apps through visual framing. In *Proc. INTERACT*.
20. Kay Connelly, Ashraf Khalil, and Yong Liu. 2007. Do I do what I say?: Observed versus stated privacy preferences. In *Human-Computer Interaction–INTERACT 2007*. Springer, 620–623.

21. Justin Cranshaw, Jonathan Mugan, and Norman Sadeh. 2011. User-controllable learning of location privacy policies with gaussian mixture models. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*.

22. EFF. 2013. Awesome Privacy Tools in Android 4.3+. (2013). https://goo.gl/atrXUB

23. William Enck, Peter Gilbert, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. 2010. TaintDroid: an information flow tracking system for real-time privacy monitoring on smartphones. *Comm. ACM* (2010).

24. Facebook. 2014. Making It Easier to Share With Who You Want. (2014). http://newsroom.fb.com/news/2014/05/making-it-easier-to-share-with-who-you-want/

25. Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A library for large linear classification. *The Journal of Machine Learning Research* 9 (2008), 1871–1874.

26. Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*. ACM, 351–360.

27. A.P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. 2012b. Android Permissions: User Attention, Comprehension, and Behavior. *Proc. of SOUPS* (2012).

28. Adrienne Porter Felt, Erika Chin, Steve Hanna, Dawn Song, and David Wagner. 2011. Android Permissions Demystified. In *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11)*. ACM, New York, NY, USA, 627–638. DOI: http://dx.doi.org/10.1145/2046707.2046779

29. Adrienne Porter Felt, Serge Egelman, and David Wagner. 2012a. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proc. SPSM*.

30. Drew Fisher, Leah Dorner, and David Wagner. 2012. Short paper: location privacy: user behavior in the field. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*. ACM, 51–56.

31. Huiqing Fu, Yulong Yang, Nileema Shingte, Janne Lindqvist, and Marco Gruteser. 2014. A Field Study of Run-Time Location Access Disclosures on Android Smartphones. *Proc. USEC* (2014).

32. Jie Gu, Yunjie Xu, Heng Xu, and Hong Ling. 2015. Interaction Effects of Contextual Cues on Privacy Concerns: The Case of Android Applications. In *Proceedings of the 2015 48th Hawaii International Conference on System Sciences (HICSS '15)*. IEEE Computer Society, Washington, DC, USA, 3498–3507. DOI: http://dx.doi.org/10.1109/HICSS.2015.421

33. Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions. In *Proc. CHI*. 10. DOI: http://dx.doi.org/10.1145/2556288.2556978

34. Peter Hornyack, Seungyeop Han, Jaeyeon Jung, Stuart Schechter, and David Wetherall. 2011. These Aren't the Droids you're Looking For: Retrofitting Android to Protect Data from Imperious Applications. In *Proc. CCS*.

35. Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. 2012. A conundrum of permissions: installing applications on an android smartphone. In *Financial Cryptography and Data Security*. Springer, 68–79.

36. Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proc. CHI*. ACM, 3393–3402.

37. Jennifer King. How Come I'm Allowing Strangers To Go Through My Phone? Smartphones and Privacy Expectations. In *Proc. SOUPS*.

38. Bart P Knijnenburg. 2014. Information Disclosure Profiles for Segmentation and Recommendation. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*.

39. Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proc. UbiComp*.

40. Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. 2014. Modeling Users' Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. In *Proc. SOUPS*.

41. Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling mobile app privacy and usability on smartphones: could user privacy profiles help?. In *Proceedings of the 23rd international conference on World wide web*. ACM, 201–212.

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

43. Mohammad Nauman, Sohail Khan, and Xinwen Zhang. 2010. Apex: Extending Android Permission Model and Enforcement with User-defined Runtime Constraints. In *Proc. CCS*.

44. Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126. DOI: http://dx.doi.org/10.1111/j.1745-6606.2006.00070.x

45. Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. In *Conference on Human factors in computing systems (CHI '03)*. ACM, New York, New York, USA, 129–136. DOI: http://dx.doi.org/10.1145/642633.642635

46. R. Parasuraman, T.B. Sheridan, and Christopher D. Wickens. 2000. 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 (May 2000), 286–297. DOI: http://dx.doi.org/10.1109/3468.844354

47. A. Patrick and Steve Kenny. 2003. From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *Proc. PET '03*. Springer. DOI: http://dx.doi.org/10.1007/978-3-540-40956-4_8

48. Anand Paturi, Patrick Gage Kelley, and Subhasish Mazumdar. 2015. Introducing Privacy Threats from Ad Libraries to Android Users Through Privacy Granules. In *Proceedings of NDSS Workshop on Usable Security (USEC'15)*. Internet Society.

49. Pew Research Center. 2014. Internet Project/GFK Privacy Panel. http://www.pewinternet.org/files/2015/05/Privacy-and-Security-Attitudes-5.19.15_Topline_FINAL.pdf. (2014).

50. J Ross Quinlan. 2014. *C4. 5: programs for machine learning*. Elsevier.

51. Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman M Sadeh. 2009. Capturing social networking privacy preferences. In *Privacy Enhancing Technologies*. Springer, 1–18.

52. Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015a. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 1–17. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub

53. F. Schaub, B. Konings, and M. Weber. 2015b. Context-Adaptive Privacy: Leveraging Context Awareness to Support Privacy Decision Making. *Pervasive Computing, IEEE* 14, 1 (Jan 2015), 34–43. DOI: http://dx.doi.org/10.1109/MPRV.2015.5

54. Scikit-Learn. 2015. Scikit-Learn Manual. http://scikit-learn.org/stable/modules/generated/sklearn.cluster.AgglomerativeClustering.html. (2015).

55. Fuming Shih, Ilaria Liccardi, and Daniel J. Weitzner. 2015. Privacy Tipping Points in Smartphones Privacy Preferences. In *Proc. CHI*. ACM.

56. Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use. In *Proc. CHI*. 10. DOI: http://dx.doi.org/10.1145/2556288.2557421

57. Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. 2014. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proc. CHI*. ACM, 91–100.

58. Richard H Thaler and Cass R Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

59. Scott Thurm and Yukari Iwatani Kane. 2010. Your apps are watching you. *The Wall Street Journal* 17 (2010).

60. Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proc. CHI*.

61. Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. 2013. Privacy manipulation and acclimation in a location sharing application. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 549–558.

62. Pamela Wisniewski, Bart P Knijnenburg, and H Richter Lipford. 2014. Profiling facebook users privacy behaviors. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*.