Towards Understanding the Influence of Personality on Mobile App Permission Settings

Frederic Raber, Antonio Krueger

DFKI, Saarland Informatics Campus, 66123 Saarbrcken
{frederic.raber,krueger}@dfki.de

Abstract. In this paper we investigate the question whether users' personalities are good predictors for privacy-related permissions they would grant to apps installed on their mobile devices. We report on results of a large online study (n=100) which reveals a significant correlation between the user's personality according to the big five personality scores, or the IUIPC questionnaire, and the app permission settings they have chosen. We used machine learning techniques to predict user privacy settings based on their personalities and consequently introduce a novel strategy that simplifies the process of granting permissions to apps.

1 Introduction

In earlier days of smartphones, users had no ability to choose the permissions each app received. Every app had a fixed set of permissions, which the user had to accept prior to installing the app. If she did not agree with even one of the requested permissions, she had no other choice than to not install the app.

This changed with Android 4.3 ("Jellybean"), where a hidden permission manager called *AppOps* for the first time gave the users the ability to change the individual permissions of an app. Since Android OS 6.0, the permission manager is integrated and visible to the user by default. For each of the permissions an app requests, the user is given a switch to either allow or deny the permission. On average, an average user has 95 apps [25] with five permissions on average [26] for each of them, resulting in a massive amount of 475 permission settings in total. According to earlier work, many users are unaware of, or at least uncomfortable with, permissions they granted to their apps[11,9,13,18].

Even worse, apps request more permissions than they need for operation[6]. Despite this fact, the average smartphone user is not aware of the risks that remain in relation to this. App stores, and also the community ratings, which are the basis for app decisions of most users, do not indicate these privacy risks at all [6]. Felt et al. [10,11] and Kelley et al. [17] have shown that the current way of displaying permissions is not clear to users and ineffective in informing them about potential risks.

Apart from better visualizing permissions and the associated privacy risks and still letting the user decide on each permission, other researchers explored the possibilities of automatically predicting and setting the app permissions using different techniques. Liu et al. succesfully trained a machine learning prediction using the settings of the four million users of the *LBE Privacy Guard app* [22], or in another approach the purpose of each permission [21], to derive a set of user profiles and assign each user the permission profile she needs.

There is evidence that a user's personality, captured by the big five personality measures[7], correlates with privacy and posting behavior, for example, on Facebook [1]: *Extraverted* users have more friends, and post more statuses and likes on Facebook. Similar results could be observed for *openness*. In contrast, more conscentious subjects are less likely to "like" a post or be a member in a large number of groups. There is also a correlation of the personality and mobile apps that are chosen by the users, and vice versa it is possible to derive the personality of a user given the installed apps on her smartphone [31]. Although it is known that personality corresponds to the usage pattern and privacy behavior on online social networks, there has not been a deeper look into the effect of personality and privacy attitudes on the choice of permission settings, and how such a correlation can be facilitated to generate an individual permission settings profile for each user.

In this paper, we try to shed light on this question and explore the influence of personality and privacy on the choice of permission settings on mobile apps, using the big five personal inventory and the IUIPC[1] questionnaire. This paper does not advance machine learning or AI, nor does it test a fully working prototype, but it bridges the gap between those two. The core contribution is the feasability analysis of the application of these techniques to mobile app privacy/security; and the derivation of design guidelines for future user interfaces in that research area. In detail we try to solve the following research questions:

1. Is there a correlation between the **general** personality measures (e.g. Big Five) and the app permission choice?
2. Is there a correlation between the **privacy attitude** (e.g. IUIPC) and the app permission choice?
3. Are there correlations inside the permission settings ?
4. Are the correlations big enough to be facilitated within a machine learning prediction?
5. How could a system look like that uses machine-learning based prediction to support the user during his privacy setting process?

We conducted a user study to capture the privacy attitudes of 100 users, together with their desired app permission settings. The results have been used as training data for a *privacy wizard*, that automatically sets the individual app permissions based on machine learning. In addition to this static approach, we examined a *Dynamic Permission Settings Prediction*, which observes the user as she adapts the Permission Settings of an app, and proposes additional changes based on the user's input on the fly. The results show that both the static as well as the dynamic setting prediction perform better than the current standard.

[1] Internet Users' Information Privacy Concerns

These two approaches allow us to support two different use-cases that capture the typical privacy setting behavior of most smartphone users: The static approach supports a use-case where a user just bought a new smartphone, and wants to set all permissions togehter. In contrast to this the dynamic approach targets for the group of users that already own a smartphone, and supports them in adapting the settings of the apps from time to time.

2 Related Work

Privacy and Personality Questionnaires There exist several approaches for measuring a person's privacy attitude within different domains. One of the earliest publications in this field is Alan Westin's work on consumer privacy indices, which was later summarized by Kumaraguru and Cranor [19]. Westin proposed three different categories of users to express their privacy attitudes: The *Unconcerned* hardly care about their privacy and tend to publish all information to the entire audience of a network. *Fundamentalists* in contrast try to disclose as little information as possible in order to preserve their privacy. The third group of persons, the *Pragmatists*, attempt to keep a balance between privacy and usability: Pragmatists believe that privacy is an important aspect, but on the other hand accept the necessity to share information in order to benefit, for example, from an additional app feature.

Although the Westin Categories have been widely used in research, the concept has several design flaws, as Woodruff and Pihur discovered recently [30]. The members of the different categories do not behave significantly differently in their actions regarding privacy. Especially the coarse-grained categorization into three categories makes it hard to predict the user's reactions to hypothetical scenarios or permission settings. The authors critisize the questionnaire as too unspecific to capture any significant effects.

The PCS[2] questionnaire [3] is more detailed and consists of 28 questions in four categories: General Caution, Technical Protection and Privacy Concern. Although more detailed, the questionnaire still adresses the general privacy attitude of a person, and not the specific context of app privacy and privacy in the context of online companies.

In contrast, the CFIP[3] [29] and the IUIPC [24] questionnaire based on it, were designed explicitely to measure the privacy of internet users, especially in the context of online shopping companies and their data collection. The authors found that the privacy attitude regarding online companies can be well expressed using three privacy measures: The *control* measure, which determines how far a subject desires to have control over the disclosure and transfer of her personal information, the desired *awareness* on how and to whom the personal information is disclosed, and *collection* describing how important it is for the subject to know which personal data is collected. As the IUIPC is the privacy questionnaire which best fits the goals of our paper, we used it in the survey of our main study.

[2] Privacy Concern Scale

[3] Scale of Concern For Information Privacy

The big five personal inventory, first created by Costa and McCrae [7], is currently the most widely accepted questionnaire to capture a person's personality. Although most reviews are very positive [15,7], there are also some critical voices [2]. Nevertheless, it is established as the standard personal inventory questionnaire. The big five is a questionnaire (also called the NEO-PI-R) in its newer form consisting of 240 items, resulting in five personality measures: *Openness to experience*, denoting general appreciation for art, emotion, adventure etc.; *Conscientiousness*, meaning the tendency to show self-discipline; *Extraversion* meaning higher or lower social engagement; *Agreeableness* in terms of cooperation with other people and *Neuroticism* as the tendency to experience negative emotions. The questionnaire in its original version is very long and requires up to 30 to 40 minutes for completion, making it unsuitable in most scenarios. Our scenario also requires a shorter solution, as we cannot prompt a user to fill in a 40-minute questionnaire before the first use of a permission recommendation app. Gosling et al. developed a shorter version to capture the big five personality traits, consisting of only ten questions [12]. Although the precision of this so-called Ten Item Personality Measure (TIPI) is not as good as with the NEO-PI-R, the results can still precisely describe the personality of a subject. The "big five" of personality can also be extracted out of written text, e.g. blog or social network entries [5]. The user burden for gathering the big five personality measures can therefore be reduced to a minimum. As stated in the introduction, there is evidence that personality correlates with the Facebook sharing behavior. We also expect some effects on the permission settings of mobile apps, and therefore included the TIPI questionnaire in our study.

Permission prediction techniques Privacy settings prediction has been a popular topic in several domains, among others online social networks. Fang and LeFevre [8] proposed a semi-supervised machine learning technique to infer privacy settings of a user's social network (SN) friends: The user is asked to label several of her friends on the SN with privacy privileges. The decision on how many and which friends have to be labeled is made by their algorithm. After this annotation phase, the software predicts the privacy privileges for the remaining, unlabeled friends.

Ravichandran et al. [28] propose the use of privacy templates for each user, in the context of location sharing with mobile apps. They observed 30 users using a mobile phone app and asked them to annotate their privacy desires towards location sharing (share location/do not share location) whenever they changed their context, e.g. when they came home from work. The app recorded the time when a context change appeared, as well as the corresponding privacy desires. Using decision trees and clustering techniques, they created several privacy profile templates. Their experiment has shown that with only three templates, the preferences of a user are matched with 90% accuracy.

There are several publications describing the prediction of mobile application settings using different data sources for the prediction. Other approaches use machine learning to predict the settings [22,21,20]. Ismail et al. [14] describe an approach which facilitates crowdsourcing in order to find an optimal tradeoff between denied permissions and usability of the app, tailored to an individual user. Liu et al. [22] use a large online database of the LBE Privacy Guard app, containing the app settings of 4.8 million users, as training data for their prediction using a linear support vector machine. 90% of the user records are used for training, 10% for testing the accuracy of the prediction. When it comes to prediction, the system uses 20% of the app settings of a user to predict the remaining 80% of settings. They used only permissions, the user and the app id for the prediction to achieve a precision score of 64.28% to 87.8%, depending on the features used. Privacy or personality attitudes were not taken into account. A similar work [21] used feedback to suggest the permission settings: For each critical permission, the system gives the user an overview on other apps and their usage frequency of the questionable permission. The user is then asked whether she feels comfortable with the previous usage or not. Based on this feedback, permission settings are recommended for the app. In total, 78.7% of the recommendations were accepted. Nevertheless, the approach needs the knowledge about the permission usage frequency of the already installed apps, and can therefore not be applied if a user just started using a new smartphone (known as the *cold start problem*). The last related work to be mentioned here [20] uses static code analysis to reconstruct the purpose of each app permission. Privacy preferences that reflect people's comfort with a permission's purpose are used to cluster the settings and to gather a finite set of privacy profiles. These profiles can later be used to assign the appropriate set of permissions for each individual user.

To conclude the related work on permission prediction, there have been several approaches using crowdsourcing or machine learning techniques like clustering, based on the permission settings themselves or using comfort with the purpose of a permission. The effect of personality on the choice of permissions has to the best of our knowledge not been explored so far. In the next sections, we will describe a system which predicts the app permission settings using the privacy attitude or personality of a user. Unlike other related work [21], our approach does not need any knowledge about previous smartphone usage behavior, and can therefore be seen as a first step towards solving the cold start problem in this scenario.

3 User Study

We conducted an online user study to discover correlations between the personality or privacy attitudes of a person, and her app permission settings. The focus was hereby on discovering *whether* there are correlations that can be used for a prediction, rather than measuring how strong the correlation could be with a large training set. To avoid side-effects, we decided to record the personality and privacy measures directly using a questionnaire rather than trying to infer the data from online social network behavior. As discussed in the last section, broader questionnaires like Westin's categories or the CFIP lead to suboptimal results. Therefore we used the more specific IUIPC questionnaire in order to capture the privacy wishes of the subjects. The personality was captured using the big five personality measure [7], more specifically the abbreviated Ten Item Personality score (TIPI) [12], which is a compressed version of the big five scale using only ten questions in total. Although possible, we did not extract the personality measures but used the TIPI for this study, to reduce any possible side-effects caused by the derivation of the measures.

In addition to these two questionnaires, we posed two additional questions regarding privacy and privacy invasion (see Table 1). In detail, we asked the subjects how recently they have been a target of a privacy invasion (five point ordinal scale from very frequently to never), and how often they enter wrong information on purpose on online websites (percentage as a numeric scale). The next subchapter describes the participants, procedure and parts of the survey in greater detail.

Label	**Question**
Falsify	Some websites ask you for personal information. When asked for such information, what percent of the time would you falsify the information?
Invasion	Have you ever been the target of a privacy invasion (e.g. your data was misused or shared without your knowledge)?

Table 1. Question text and label of the additional question set.

3.1 Methodology

The study was conducted as an online survey using the software LimeSurvey[4]. 100 participants were recruited using Prolific Academic,[5] which allows us to select only active Android smartphone users with at least three of their own apps installed. Users had to conduct the study at a PC or Laptop at home. Studies in the past have shown that participants who are recruited via online services, like in our case, lead to a similar quality of the results, like participants recruited at a university [4]. The participants were paid a compensation of 2£ upon successful participation. To motivate the subjects to fill out the questionnaire honestly, the compensation was only paid after the submitted data was checked for plausibility by us. If the result of a subject was rejected, for example if she failed to answer the control questions correctly, a new participant has been recruited to fill in the gap. Therefore we have exactly 100 viable results.

[4] `https://www.limesurvey.org`, last accessed 09-05-2016

[5] `https://www.prolific.ac/`, last accessed 09-05-2016

The age of the participants ranged from 18 to 61 years (average 30.13, SD 8.53). The recruited audience was very diverse: We recruited students, self-employed workers, employees, and also homemakers.

The survey can be divided into two parts: In the first part, we asked the subjects to fill out the above described privacy and personality questionnaires. In the second phase, we asked them to look up the permissions of their *up to ten* most frequently used apps, and let them enter them into the survey form (see upper right of Figure 1). Next to each of the permissions, just like in the Andoid OS 6.0 interface where permissions are allowed per default, we asked them to state whether they would reject the permission if they could. The second column ("I would revoke the permission") is only active for permissions that are marked as "owned by the app" in the first column. According to previous work[13], users hardly know which permissions are requested and how they can determine which ones are used. Therefore the subjects were given simple step-by-step instructions including screenshots of every step (see lower right of Figure 1), in order to make sure they are able to retrieve the permissions for their apps correctly. To make sure users can conduct the task correctly, the questionnaire asked to enter the permissions of a specific app (namely google maps). Only if the task was done correctly, participants were allowed to continue. Different app versions or different android os versions can request a different set of permissions, therefore we only checked whether different subjects entered a different permission set for the same application, if the version was the same. In our study, this was not the case.

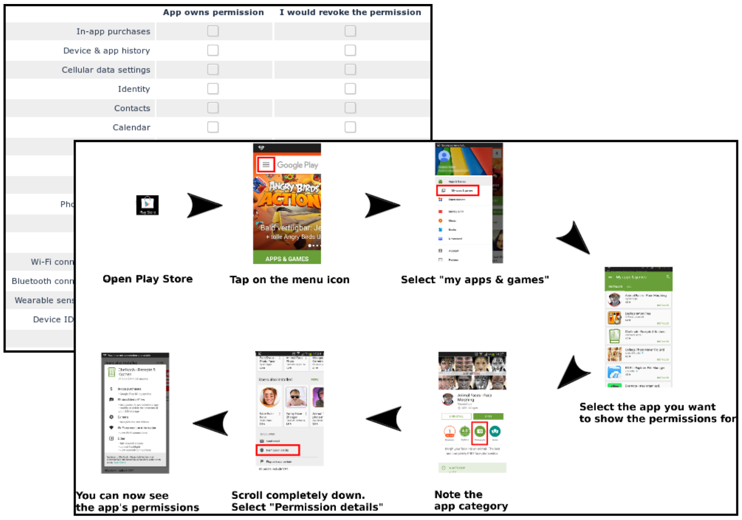

	App owns permission	I would revoke the permission
In-app purchases	☐	☐
Device & app history	☐	☐
Cellular data settings	☐	☐
Identity	☐	☐
Contacts	☐	☐
Calendar	☐	☐
Pho		
Wi-Fi conn		
Bluetooth conn		
Wearable sens		
Device ID		

Open Play Store
Tap on the menu icon
Select "my apps & games"
Select the app you want to show the permissions for
Note the app category
Scroll completely down. Select "Permission details"
You can now see the app's permissions
Google Play
APPS & GAMES

Fig. 1. Step-by-step instructions for permission retrieval given to the subjects (lower right) and one of the ten questionnaire pages for capturing the app permissions and settings preference (allow/reject) in the upper left.

The survey ended with a short feedback question in free-text style.

3.2 Results

The 100 participants entered in total the settings of 876 apps into the system. On average each user filled in the details of three to ten apps, 8.65 on average. Figure 2 shows a detailed graph on the number of settings with a specific amount of denied permissions. In most cases (447 out of the 876 settings), no permission was denied. The answers to the different items of the IUIPC and TIPI questionnaires have been reversed if needed, and combined to the according three (IUIPC) or five (TIPI) personality measures, as described in literature[24,12]. We only used these combined measures for the machine learning, as well as for the statistical analysis.

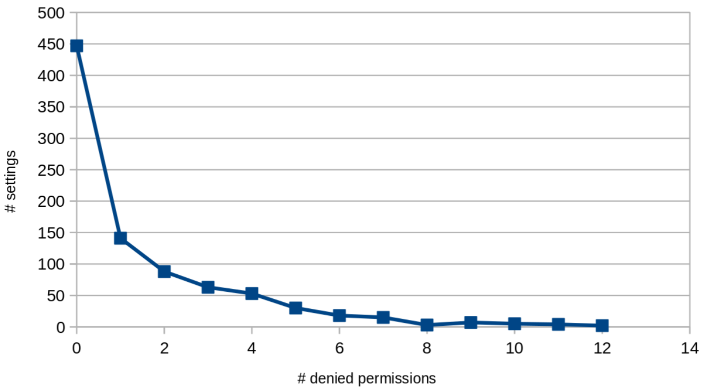

Fig. 2. Number of settings with a specific number of denied permissions.

Table 2 shows how often each of the permissions was denied throughout the study.

For each participant and app permission, we computed a *permission coefficient*, that denotes how often the permission is denied by this user. The permission coefficient $comb$ is computed as $comb = \frac{|rejected|}{|used|}$, where $|used|$ denotes how many of his four to ten apps used the permission, and $|rejected|$ how often the permission has been rejected. For each participant, we thereby have 17 *permission coefficients*, one for each permission, that give us the normalized likelihood of a permission to be denied. These coefficient values range continuously from 0 (never denied) to 1 (always denied) and are independent between participants. In the next step, we wanted to find out whether the questionnaire answers correlate with the *permission coefficients*. According to the shape of our data (independent, mostly ordinal values, not necessarily normal-distributed),

Permission	% denied
Purchase	18.4
History	17.6
Cellular	9.5
Identity	26.6
Contacts	36.3
Calendar	16.7
Location	34.6
SMS	37.2
Phone	28.6
Photos	30.9
Camera	28.9
Microphone	25.1
Wifi	12.8
Bluetooth	5.0
ID	31.0
Other	17.9

Table 2. Percentages of denies for each app permission.

we decided to use a non-parametric test, and therefore performed a Spearman correlation ("Spearman's Rho") on the results of the questionnaire and the permission coefficients. The results are shown in Tables 3 and 4. The measures of the privacy and personality questionnaires are on the rows, whereas the app permission coefficients are plotted as the columns of the table. Significant and highly significant correlations are marked with one or two asterisks, and colored in gray or dark gray, respectively. Note that for our purpose, the correlation coefficient is more important than the significance, as it denotes the ascent of the regression line between the data points. The higher the value, the easier it is to forecast a *permission coefficient* given the questionnaire answers. The measures of the specialized IUIPC privacy questionnaire (collection, control, awareness) received the best correlation scores, the collection measure yields high correlation coefficients which are mostly (highly) significant for most of the permissions. Control and awareness both correlate on *some* of the permissions.

The general privacy questionnaires (TIPI) received lower but still useful correlation scores. The *open to experiences* of the subjects correlates with two of the app permissions and is, together with *conscientousness* and *emotional stability* (one significant correlation each) the most expressive personality measure. Although only four personality - permission pair lead to significant correlations, the *correlation coefficients* still remain medium high for several other combinations, making it a promising candidate for machine-learning based prediction. The additional questions on the other hand, received only small correlation scores. We therefore dropped the additional questionnaires for the machine learning and the evaluation.

In addition to the correlations between privacy, personality and the permission settings, we also observed the pairwise correlations between the permission settings. In contrast to the correlation between questionnaire and permissions, not only some but the whole lot of pairwise correlations are highly significant with correlation coefficients reaching from 0.217 (Phone - Wifi) up to 0.859. The highest correlations can be achieved with the permission pairs *cellular information - identity* ($r = 0.859$), *cellular information - purchase* ($r = 0.81$) and *cellular information - history* ($r = 0.778$).

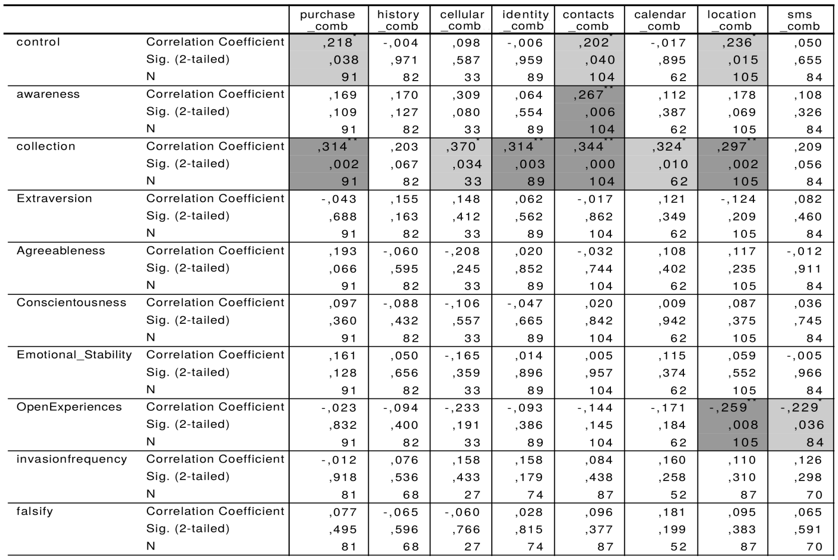

		purchase _comb	history _comb	cellular _comb	identity _comb	contacts _comb	calendar _comb	location _comb	sms _comb
control	Correlation Coefficient	,218*	-,004	,098	-,006	,202*	-,017	,236*	,050
	Sig. (2-tailed)	,038	,971	,587	,959	,040	,895	,015	,655
	N	91	82	33	89	104	62	105	84
awareness	Correlation Coefficient	,169	,170	,309	,064	,267**	,112	,178	,108
	Sig. (2-tailed)	,109	,127	,080	,554	,006	,387	,069	,326
	N	91	82	33	89	104	62	105	84
collection	Correlation Coefficient	,314**	,203	,370*	,314**	,344**	,324*	,297**	,209
	Sig. (2-tailed)	,002	,067	,034	,003	,000	,010	,002	,056
	N	91	82	33	89	104	62	105	84
Extraversion	Correlation Coefficient	-,043	,155	,148	,062	-,017	,121	-,124	,082
	Sig. (2-tailed)	,688	,163	,412	,562	,862	,349	,209	,460
	N	91	82	33	89	104	62	105	84
Agreeableness	Correlation Coefficient	,193	-,060	-,208	,020	-,032	,108	,117	-,012
	Sig. (2-tailed)	,066	,595	,245	,852	,744	,402	,235	,911
	N	91	82	33	89	104	62	105	84
Conscientousness	Correlation Coefficient	,097	-,088	-,106	-,047	,020	,009	,087	,036
	Sig. (2-tailed)	,360	,432	,557	,665	,842	,942	,375	,745
	N	91	82	33	89	104	62	105	84
Emotional_Stability	Correlation Coefficient	,161	,050	-,165	,014	,005	,115	,059	-,005
	Sig. (2-tailed)	,128	,656	,359	,896	,957	,374	,552	,966
	N	91	82	33	89	104	62	105	84
OpenExperiences	Correlation Coefficient	-,023	-,094	-,233	-,093	-,144	-,171	-,259**	-,229*
	Sig. (2-tailed)	,832	,400	,191	,386	,145	,184	,008	,036
	N	91	82	33	89	104	62	105	84
invasionfrequency	Correlation Coefficient	-,012	,076	,158	,158	,084	,160	,110	,126
	Sig. (2-tailed)	,918	,536	,433	,179	,438	,258	,310	,298
	N	81	68	27	74	87	52	87	70
falsify	Correlation Coefficient	,077	-,065	-,060	,028	,096	,181	,095	,065
	Sig. (2-tailed)	,495	,596	,766	,815	,377	,199	,383	,591
	N	81	68	27	74	87	52	87	70

Fig. 3. Correlations between the privacy/personality questions and app permission settings.

4 Permission Wizard

Based on the results of the user study, we decided to use the results as training data to predict the privacy settings of a user's apps, based on her personality and privacy attitudes. The current Android interface allows all permissions, therefore the interesting cases are the ones where at least one permission is denied. We followed the example of earlier work [21,22] and concentrated on these harder cases for our evaluation, and took only them into account for the prediction and evaluation.

Similar publications [22] used a simple SVM algorithm for their prediction. We also tried out SVM and several other classification methods, and achieved the best results with a KNeighbors implementation with two as the number of neighbors.

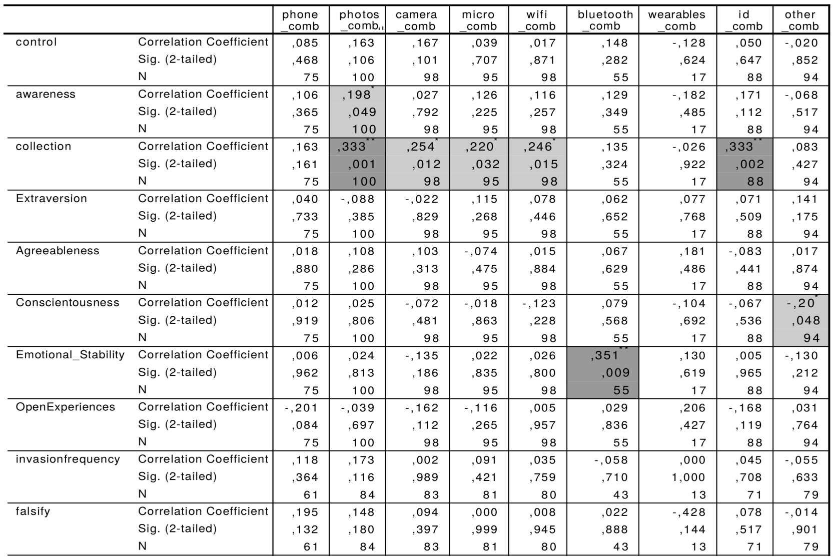

		phone _comb	photos _comb	camera _comb	micro _comb	wifi _comb	bluetooth _comb	wearables _comb	id _comb	other _comb
control	Correlation Coefficient	,085	,163	,167	,039	,017	,148	-,128	,050	-,020
	Sig. (2-tailed)	,468	,106	,101	,707	,871	,282	,624	,647	,852
	N	75	100	98	95	98	55	17	88	94
awareness	Correlation Coefficient	,106	,198*	,027	,126	,116	,129	-,182	,171	-,068
	Sig. (2-tailed)	,365	,049	,792	,225	,257	,349	,485	,112	,517
	N	75	100	98	95	98	55	17	88	94
collection	Correlation Coefficient	,163	,333**	,254*	,220*	,246*	,135	-,026	,333**	,083
	Sig. (2-tailed)	,161	,001	,012	,032	,015	,324	,922	,002	,427
	N	75	100	98	95	98	55	17	88	94
Extraversion	Correlation Coefficient	,040	-,088	-,022	,115	,078	,062	,077	,071	,141
	Sig. (2-tailed)	,733	,385	,829	,268	,446	,652	,768	,509	,175
	N	75	100	98	95	98	55	17	88	94
Agreeableness	Correlation Coefficient	,018	,108	,103	-,074	,015	,067	,181	-,083	,017
	Sig. (2-tailed)	,880	,286	,313	,475	,884	,629	,486	,441	,874
	N	75	100	98	95	98	55	17	88	94
Conscientousness	Correlation Coefficient	,012	,025	-,072	-,018	-,123	,079	-,104	-,067	-,20*
	Sig. (2-tailed)	,919	,806	,481	,863	,228	,568	,692	,536	,048
	N	75	100	98	95	98	55	17	88	94
Emotional_Stability	Correlation Coefficient	,006	,024	-,135	,022	,026	,351**	,130	,005	-,130
	Sig. (2-tailed)	,962	,813	,186	,835	,800	,009	,619	,965	,212
	N	75	100	98	95	98	55	17	88	94
OpenExperiences	Correlation Coefficient	-,201	-,039	-,162	-,116	,005	,029	,206	-,168	,031
	Sig. (2-tailed)	,084	,697	,112	,265	,957	,836	,427	,119	,764
	N	75	100	98	95	98	55	17	88	94
invasionfrequency	Correlation Coefficient	,118	,173	,002	,091	,035	-,058	,000	,045	-,055
	Sig. (2-tailed)	,364	,116	,989	,421	,759	,710	1,000	,708	,633
	N	61	84	83	81	80	43	13	71	79
falsify	Correlation Coefficient	,195	,148	,094	,000	,008	,022	-,428	,078	-,014
	Sig. (2-tailed)	,132	,180	,397	,999	,945	,888	,144	,517	,901
	N	61	84	83	81	80	43	13	71	79

Fig. 4. Correlations between the privacy/personality questions and app permission settings continued.

As we only have $2^9 = 512$ combinations of input features, we were able to determine the optimum set of features by using a brute-force procedure of training and selecting the features according to the precision of each of the combinations one after another. We followed the usual way for training, adjusting parameters, and validating the prediction of a machine learning algorithm. We used a ten-fold cross validation to prevent a biasing of the data. In this validation method, the data is split into ten parts of the same size, and the validation procedure is performed ten times: In each of the ten runs, the data set is split into two basic parts: The first part is called the *training set*, and is composed of 90% of the data set. It is used to train, to calibrate the prediction algorithm and select the optimal features. The second and remaining part is called the *test set*, and is **not** used for training and fitting, it remains untouched. It is used for the evaluation of the results later. We used 80% of the *training set* (**not** the test set, as this remains untouched until the evaluation in the next chapter) to fit the algorithm, and 20% to find the optimal set of features using the above-mentioned brute-force method. After each run, another of the ten splits is used as the *test set*, and the remaining splits for the *training set*. After the ten distinct runs, we used the average precision of all runs for selecting the best set of features. Table 3 shows the features selected for the prediction using the IUIPC or the big five personality measures as input features.

The selected features correspond to the measures with the highest correlation in Tables 3 and 4, supporting the correctness of the selection method.

Feature Set	Selected features
IUIPC	Collection, Control
Personality	Extraversion, OpenExperiences

Table 3. Selected features for each of the two feature sets IUIPC and personality.

Evaluation We evaluated the results of the prediction following a similar approach as for the selection of input features using ten-fold cross validation: First the prediction is trained with the *training set* consisting of 90% of the data. This time, we used a static set of input features as described in Table 3, that was not changed throughout the evaluation procedure. Afterwards, the feature values from the *test set* are used to predict the permission settings, and compared with the actual permission settings of the *test set*. Again this procedure was repeated ten times, and the results were averaged.

In order to get an impression of the quality of the results, we implemented a naive approach to predict the settings, which will later be called the *baseline* or *random* condition. We started with a simple random method, which randomly predicts "allow" or "deny" for each of the permissions, giving a 50% accuracy. Since the percentage of allow and deny differs from permission to permission and is rarely at 50% for both (see Table 2), we enhanced the random approach by a probabilistic component: We first use the *training set* to calculate the probability of getting allowed or denied for each permission respectively. Based on these probabilities, we then predict the permission settings on the *test set*. If for example a setting for *Contacts* permission has to be chosen, the prediction will decide to "allow" with a probability of 63.7%, and to "deny" in 36.3% of all cases.

As before, ten runs have been conducted to evaluate the probabilistic random method, and the results have been averaged.

The percentage of correct predictions of this probabilistic *Random* approach, as well as the correctness using only the IUIPC or the personality metrics as features, is shown in Table 4. The columns denote the feature sets, whereas the rows contain the different app permissions. The topmost row ("all") denotes the average percentage over all permissions.

Although the probabilistic approach achieves significantly better results ($M = 59,64$) than a pure random method, the machine learning-based prediction can still outperform it with both feature sets (M_{IUIPC} = 70.92, $M_{Personality}$ = 69.37). Best results can be achieved for the bluetooth (M_{IUIPC} = 96.66, $M_{Personality}$ = 93.33) and cellular info permissions (M_{IUIPC} = 92.5, $M_{Personality}$ = 91.25). The location permission was hardest to predict (M_{IUIPC} = 53.33, $M_{Personality}$ = 58.48). Overall, the machine learning approach outperformed the random probabilistic method by more than 10%.

Permission	Random	IUIPC	Personality
All	59.64	70.92	69.37
Purchase	59.37	78.13	67.50
History	65.88	72.94	78.82
Cellular	78.75	92.50	91.25
Identity	51.87	68.44	60.62
Contacts	48.88	55.18	64.44
Calendar	70.00	80.00	81.11
Location	45.15	53.33	58.48
SMS	54.37	50.00	57.50
Phone	53.33	67.33	58.66
Photos	47.31	63.65	62.44
Camera	53.92	60.00	61.07
Microphone	52.50	74.00	69.00
Wifi	68.82	86.47	78.82
Bluetooth	84.44	96.66	93.33
ID	56.08	64.78	58.70
Other	63.55	71.33	68.22

Table 4. Prediction accuracy (in percent of correct predictions) for the prediction with the Random Probabilistic Model (Random), and prediction using the IUIPC questionnaire or the Big Five Personality test.

4.1 Dynamic Setting Prediction

Besides the prediction of all settings at once using the personality and privacy measures, we discovered techniques for how the pairwise correlation between the permission settings can be used to actively support the smartphone user during her decision process, while setting the permission settings. Given that the permission settings are displayed as a scrollable list as on Android OS, we assume that most smartphone users traverse the list from the top to the the bottom, and change the permissions they want to set to "deny". As soon as a change is made, we can take this change as well as the permission settings above this entry, as an input to predict the remaining settings below. This technique will later be called *dynamic setting prediction*. In detail, we trained estimators for all possible compinations of selected permissions, and serialized them to a file which is loaded into a cache at startup of the application. As soon as the user interacts with the settings, the estimator corresponding to the selected settings is retrieved from the cache and used for the prediction.

We used the study data to simulate the subject's behavior when setting the permission settings of all of her apps, either without support, or with the support of the *dynamic setting prediction*. We observed the prediction accuracy using only the already-set permissions, as well as the set permissions in addition to the IUIPC, personality, and all together. The distribution of study data to training and test set is the same as described in the last chapter: 90% for training, and 10% as the *test set* for validating the correctness of the prediction.

The procedure for the validation works as follows: Initially, the predicted setting allows all permissions. Then we traverse each permission of the setting to be evaluated, one after another, and check whether the prediction meets the actual permission setting. If not, the predicted setting is adjusted, and the remaining permissions below are predicted based on the permissions above. Whenever this occurs, a user interaction (later called a *click*) is recorded. The validation in pseudo-code is described below.

```
# traverse all the user settings
for each user_settg in testset:

 # initially, all settings are
 # set to "allow"
 pred=allow_all

 for each perm in user_settg:
  if user_settg[perm]!=pred[perm]:

   # prediction was wrong,
   # user had to change the setting
   #-> predict remaining settings

   pred[perm]=user_settg[perm]
   predict_settings_below()
```

We compared the number of *clicks* needed when using the *dynamic setting prediction* to a case where the user simply clicks on all permissions she wants to deny, without any support, as it is currently implemented on Android. As our prediction technique requires a user input, e.g. setting at least one permission to "deny", we used only app settings for the evaluation where one or more permissions were denied.

Table 5 shows the results of the evaluation procedure. For each of the input sets described above, we compared the average clicks needed for each app setting with ("Clicks (supported)") and without ("Clicks (unsupported)") the support of the Dynamic Setting Prediction. Columns one to three describe the percentage of cases where the user needed fewer clicks ("won"), the same amount of clicks ("draw"), or more clicks ("lost") with the prediction than without any support.

The prediction works best when using all features, e.g. the previously set permissions, IUIPC, Personality and Additional questions as described in Table 3. In that case, the rate of needed clicks for the *dynamic setting prediction* drops to an average of 1.58 per setting, compared to 2.00 for the unsupported case. 91.89% of the settings require at maximum the same amount of clicks, 24.66% even less clicks than the unsupported version. Only in 8.11% of the cases did the user need more interactions with the prediction enabled. The ratio of needed clicks and the win/lose ratio slightly decreases with a decreasing feature set.

Input	Won %	Draw %	Lost %	Clicks(supp.)	Clicks(unsupp.)
Only Permissions	23.49	59.40	17.15	1.91	2.22
IUIPC	25.76	60.60	13.63	1.83	2.21
Personality	26.58	59.30	14.12	1.70	2.10
All	24.66	67.23	8.11	1.58	2.00

Table 5. Results of the dynamic settings prediction, using only the previously selected permissions, or the permissions in addition to the IUIPC questionnaire, the Big Five Personality Score, our additional questionnaire or all previously mentioned questionnaires together.

Using only the previous permissions as the prediction input, the *dynamic setting prediction* needed on average 1.91 clicks for each setting, compared to 2.22 clicks without the prediction. In 82.89% of the cases, the prediction needed the same or fewer clicks, whereas in 17.15% of the settings, the unsupported version required less user interaction.

5 Discussion and Limitations

We were able to prove significant correlations both between the general personality (captured by the TIPI questionnaire) as well as the privacy attitude and the app permission choice. Furthermore, the correlations are powerful enough to train a machine learning algorithm, that is able to predict these settings based on the personality, with a precision of more than 70 %. The machine learning can be used in two different use-cases, first in a traditional privacy wizard, and second in a *dynamic* approach that supports the user on the fly while she is doing her settings, as described below. Nevertheless, there are still some points which can be improved, as discussed in the following subsections.

5.1 Possible implementations of the approaches

The two techniques presented in the former chapters can be used to implement two different use-cases: In the first use-case, a smartphone user buys a new smartphone, and has to enter his app permission settings for the first time. This is also the case, if the older smartphone was running an Adnroid version below 6.0, where app permissions are not supported. As mentioned in the related work section, users are either overchallenged by the technicality and complexity of the app permissions, or fear the burden of setting every single permission for each app they use. A privacy wizard based our machine learning estimators reduce this problem: An implementation of our approach offers the user to either answer the twelve questions of the IUIPC questionnaire, or to connect to facebook / twitter to read in the user's posts and to extract the big five personality measures out of her written text [5]. Afterwards the wizard suggests a complete set of permission settings, that can be reviewed by the user. Especially lay users profit from the questionnaire, as it contains only non-technical questions that are easy to answer.

Furthermore, both user groups save a lot of time by answering only twelve IUIPC questions (or by just connecting to facebook) instead of on average 475 distinct app permission settings.

The second use-case considers a user which has a running smartphone, who wants to individually set the permission settings rather than to trust a permission wizard. In this case, the dynamic settings prediction could be integrated into Android's permission setting dialogue. When the users traverses the list of app permissions from top to the bottom (which can sometimes take a while), Android recognizes when a permission is set to deny, and denies also other permissions further below, which might also be denied according to our dynamic privacy prediction. Changes are marked in orange, so that the user can see which changes have been made automatically, and review them. According to the study results, this technique should save the user a significant amount of interaction (clicks) and save time which should also lead to a smaller frustration.

5.2 Different precision for different questionnaires

Both the IUIPC, as well as the personality questonnaire performed well in the validation study, giving an average precision of 70% and 69% of correct predictions. Although the personality measures can be automatically extracted in contrast to the IUIPC, we do not recommend to stick only to that questionnaire: Having a closer look at Table 4 reveals that both the IUIPC and the personality questionnaire complement each other: Permissions which are hard to predict in the additional questionnaire (like Contacts, Location, SMS, History) can be better predicted using the personality measures, and vice versa. If the best questionnaire is selected for each permission, a precision of 72.80% can be achieved within our test data. Whether that assumption holds for larger data sets has to be proven in future work.

5.3 Limited size of the training data set and precision of the prediction

The results in Figures 3 and 4 indicate there is a strong correlation between the answers to the personality and privacy questionnaires, and the permission settings of the apps. We were also able to predict these settings using machine learning. The proposed approach predicted about ten percent more of the permission settings correctly compared to the naive approach. Compared to related work like Liu et al. [22,20], we did not have the possibility to draw on a large online database, containing millions of datasets. As we need the personality measures in addition to the permission settings, we had to gather the training data in an online survey, and therefore have, compared to the mentioned work, a relatively small dataset. Thus the performance of our personality-based prediction cannot directly be compared to these approaches. As usual with a machine learning approach, we expect the prediction precision to increase with an increasing data set size. We would like to explore how our approach performs with a large training database in the future, and compare it to other systems that base their prediction on the permissions of similar users[20], or the purpose of the permission[22].

5.4 Control of context factors

Each experiment faces the problem of contextual factors that cannot be controlled, we designed the experiment in a way to decrease these factors to a minimum. By using prescreening, our survey could only be edited using a laptop or computer, and not on the go using a mobile phone. We further required the participants to be at home. We were therefore able to assure the same location/occasion for each participant. There are still some context factors left, like general distrust towards app producers, that cannot be controlled. Similar work also discovered that the purpose of the permission has an impact on the decision [20], which we would also like to add in a future version. To avoid personality biasing, we compared the results of the personality questionnaire with the mean values that have been recorded in earlier research. With our final data set, we could not prove any significant difference with one of the five personality traits.

5.5 Number of denied permissions per setting and dynamic prediction accuracy

The dynamic prediction of permission settings often has the same amount of clicks as the unsupported permission setting procedure, ranging from 59.4% without additional features up to 67.23% with all features enabled. The dynamic prediction clearly profits from user interaction; the more permissions the user sets to deny, the more input features are available for the prediction. Having a look at Figure 2, we can see that about one third of the permissions that have at least one denied permission, have *exactly* one denied permission. In these 33% of all cases, the dynamic prediction does not give any advantage, as it starts predicting only after the first user interaction. To the contrary, it is even possible that the algorithm predicts one of the following permissions as “deny”, leading to more clicks than without the prediction. An additional 20% have only two denied permissions, which are also hard for the approach to predict. Despite these difficulties, the prediction still needs fewer clicks on average than the unsupported approach.

5.6 Prediction precision for the different permissions

The prediction accuracy greatly differs between the different permissions that Android offers to the apps. Comparing Tables 2 and 4, we can see that these differences negatively correlate with the percentage of denial of a permission. When a permission is often allowed, the prediction accuracy, especially for the probabilistic model, is also high. Therefore the *location information*, as well as the *SMS* and *contacts* permissions are hardest to predict with the probabilistic model. On the other hand, these are the categories where the prediction outperforms the random probabilistic model the most.

5.7 Future work

We took a first step towards an automatic permission settings prediction using personality and privacy attitude as input features. We could prove a significant correlation between those two, and were able to design a static, as well as an on-the-fly prediction for the permissions. Nevertheless, the evaluation was done only theoretically, using the study data. In future work, we would like to implement the described approach as a mobile app, and conduct a lab study on the effects of the prediction support. We are especially interested in how far the dynamic prediction will be accepted. Although in our theoretical study, the approach works as well as or better than a standard android permission manager in more than 91% of the cases, we would like to confirm our results in an applied scenario, where people interact with concrete examples rather than hypothetical ones.

In a second step, an in-the wild study would be desirable where a prototypical app is released to the app store. Having a larger user base and additional training data would hopefully help to improve the prediction accuracy, and make it possible to compare the prediction to related work with a similar user base[22,20].

The decision for location privacy settings in Social Networks depends also on context factors like the current location of the user, the occasion, or the purpose for a location retrieval [27,23]. The context a permission is given might also have an effect on the choice of mobile app permission settings, for example if a user wants to grant the web browser access to the internal storage only to store e-ticket receipts. The domain of app permissions further has several contexts which require a temporary or one-time permission grant, for instance when an app needs access to the SMS permission for sending a registration SMS once [16]. Context information could be easily integrated into our prediction system as an additional feature.

6 Conclusion

The Android permission system is powerful, but the maintenance of each of the apps' permissions is very cumbersome and time-consuming. Related work already elaborated on the prediction of these settings using large online permission setting databases, crowdsourcing approaches or privacy profiles based on the permission purpose. We conducted an online user study with 100 participants to discover the effects of personality and privacy attitude on the permission settings. We found strong correlations between personality, privacy attitudes and the settings. We evaluated two approaches for supporting the user during her decision process. The first uses the personality and privacy measures to directly predict all app settings, whereas the second supports the user while doing the permission settings, taking the user's interaction as an input to predict the settings not done so far. Although the training set is very small compared to related work, we were able to outperform the current standard with both approaches.

References

1. Bachrach, Y., Kosinski, M., Graepel, T., Kohli, P., Stillwell, D.: Personality and patterns of facebook usage. In: Proceedings of the 4th Annual ACM Web Science Conference. pp. 24–32. WebSci '12, ACM, New York, NY, USA (2012), http://doi.acm.org/10.1145/2380718.2380722
2. Block, J.: A Contrarian View of the Five-Factor Approach to Personality Description. Psychological Bulletin 117, 187–215 (1995)
3. Buchanan, T., Paine, C., Joinson, A.N., Reips, U.D.: Development of measures of online privacy concern and protection for use on the internet. Journal of the American Society for Information Science and Technology 58(2), 157–165 (2007), http://dx.doi.org/10.1002/asi.20459
4. Buhrmester, M., Kwang, T., Gosling, S.: Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? Perspectives on Psychological Science 6(1), 3–5 (2011)
5. Chen, J., Haber, E., Kang, R., Hsieh, G., Mahmud, J.: Making use of derived personality: The case of social media ad targeting. In: International AAAI Conference on Web and Social Media (2015), http://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10508
6. Chia, P.H., Yamamoto, Y., Asokan, N.: Is this app safe?: A large scale study on application permissions and risk signals. In: Proceedings of the 21st International Conference on World Wide Web. pp. 311–320. WWW '12, ACM, New York, NY, USA (2012), http://doi.acm.org/10.1145/2187836.2187879
7. Costa, P., McCrae, R., Psychological Assessment Resources, I.: Revised NEO Personality Inventory (NEO PI-R) and NEO Five-Factor Inventory (NEO-FFI). Psychological Assessment Resources (1992), https://books.google.co.in/books?id=mp3zNwAACAAJ
8. Fang, L., LeFevre, K.: Privacy wizards for social networking sites. In: Proceedings of the 19th International Conference on World Wide Web. pp. 351–360. WWW '10, ACM, New York, NY, USA (2010), http://doi.acm.org/10.1145/1772690.1772727
9. Felt, A.P., Chin, E., Hanna, S., Song, D., Wagner, D.: Android permissions demystified. In: Proceedings of the 18th ACM Conference on Computer and Communications Security. pp. 627–638. CCS '11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/2046707.2046779
10. Felt, A.P., Greenwood, K., Wagner, D.: The effectiveness of application permissions. In: Proceedings of the 2nd USENIX Conference on Web Application Development. pp. 7–7. WebApps'11, USENIX Association, Berkeley, CA, USA (2011), http://dl.acm.org/citation.cfm?id=2002168.2002175
11. Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E., Wagner, D.: Android permissions: User attention, comprehension, and behavior. In: Proceedings of the Eighth Symposium on Usable Privacy and Security. pp. 3:1–3:14. SOUPS '12, ACM, New York, NY, USA (2012), http://doi.acm.org/10.1145/2335356.2335360
12. Gosling, S.D., Rentfrow, P.J., Swann, W.B.: A Very Brief Measure of the Big-Five Personality Domains. Journal of Research in Personality 37(6), 504–528 (December 2003), http://dx.doi.org/10.1016/S0092-6566(03)00046-1
13. Harbach, M., Hettig, M., Weber, S., Smith, M.: Using personal examples to improve risk communication for security & privacy decisions. In: Proceedings of the 32nhistoryd Annual ACM Conference on Human Factors in Computing Systems. pp. 2647–2656. CHI '14, ACM, New York, NY, USA (2014), http://doi.acm.org/10.1145/2556288.2556978

14. Ismail, Q., Ahmed, T., Kapadia, A., Reiter, M.K.: Crowdsourced exploration of security configurations. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 467–476. CHI '15, ACM, New York, NY, USA (2015), http://doi.acm.org/10.1145/2702123.2702370
15. John, O.P., Srivastava, S.: The Big Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives. In: Pervin, L.A., John, O.P. (eds.) Handbook of Personality: Theory and Research, pp. 102–138. Guilford Press, New York, second edn. (1999), http://darkwing.uoregon.edu/~{}sanjay/pubs/bigfive.pdf
16. Jung, J., Han, S., Wetherall, D.: Short paper: Enhancing mobile application permissions with runtime feedback and constraints. In: Proceedings of the Second ACM Workshop on Security and Privacy in Smartphones and Mobile Devices. pp. 45–50. SPSM '12, ACM, New York, NY, USA (2012), http://doi.acm.org/10.1145/2381934.2381944
17. Kelley, P.G., Consolvo, S., Cranor, L.F., Jung, J., Sadeh, N., Wetherall, D.: A conundrum of permissions: Installing applications on an android smartphone. In: Proceedings of the 16th International Conference on Financial Cryptography and Data Security. pp. 68–79. FC'12, Springer-Verlag, Berlin, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-34638-5_6
18. Kelley, P.G., Cranor, L.F., Sadeh, N.: Privacy as part of the app decision-making process. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 3393–3402. CHI '13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2470654.2466466
19. Kumaraguru, P., Cranor, L.F.: Privacy indexes: A survey of westin's studies. ISRI Technical Report (2005)
20. Lin, J., Liu, B., Sadeh, N., Hong, J.I.: Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In: Symposium On Usable Privacy and Security (SOUPS 2014). pp. 199–212. USENIX Association, Menlo Park, CA (Jul 2014), https://www.usenix.org/conference/soups2014/proceedings/presentation/lin
21. Liu, B., Andersen, M.S., Schaub, F., Almuhimedi, H., Zhang, S.A., Sadeh, N., Agarwal, Y., Acquisti, A.: Follow my recommendations: A personalized privacy assistant for mobile app permissions. In: Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). pp. 27–41. USENIX Association, Denver, CO (Jun 2016), https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu
22. Liu, B., Lin, J., Sadeh, N.: Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In: Proceedings of the 23rd International Conference on World Wide Web. pp. 201–212. WWW '14, ACM, New York, NY, USA (2014), http://doi.acm.org/10.1145/2566486.2568035
23. Lugano, G., Saariluoma, P.: To Share or Not to Share: Supporting the User Decision in Mobile Social Software Applications, pp. 440–444. Springer Berlin Heidelberg, Berlin, Heidelberg (2007), http://dx.doi.org/10.1007/978-3-540-73078-1_61
24. Malhotra, N.K., Kim, S.S., Agarwal, J.: Internet users' information privacy concerns (iuipc): The construct, the scale, and a causal model. Info. Sys. Research 15(4), 336–355 (Dec 2004), http://dx.doi.org/10.1287/isre.1040.0032
25. Olmstead, K., Atkinson, M.: The next web. android users have an average of 95 apps installed on their phones, according to yahoo aviate data. http://www.pewinternet.org/2015/11/10/an-analysis-of-android-app-permissions/ (2015), accessed: 2016-02-01

26. Olmstead, K., Atkinson, M.: 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
27. Patil, S., Le Gall, Y., Lee, A.J., Kapadia, A.: My Privacy Policy: Exploring End-user Specification of Free-form Location Access Rules, pp. 86–97. Springer Berlin Heidelberg, Berlin, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-34638-5_8
28. Ravichandran, R., Benisch, M., Kelley, P.G., Sadeh, N.: Capturing social networking privacy preferences: Can default policies help alleviate tradeoffs between expressiveness and user burden? In: Proceedings of the 5th Symposium on Usable Privacy and Security. pp. 47:1–47:1. SOUPS '09, ACM, New York, NY, USA (2009), http://doi.acm.org/10.1145/1572532.1572587
29. Smith, H.J., Milberg, S.J.: Information privacy: Measuring individuals' concerns about organizational practices. MIS Q. 20(2), 167–196 (Jun 1996), http://dx.doi.org/10.2307/249477
30. Woodruff, A., Pihur, V., Acquisti, A., Consolvo, S., Schmidt, L., Brandimarte, L.: Would a privacy fundamentalist sell their dna for $1000... if nothing bad happened thereafter? a study of the westin categories, behavior intentions, and consequences. In: Proceedings of the Tenth Symposium on Usable Privacy and Security (SOUPS). ACM, ACM, New York, NY (2014), https://www.usenix.org/conference/soups2014/proceedings/presentation/woodruff, iAPP SOUPS Privacy Award Winner
31. Xu, R., Frey, R.M., Vuckovac, D., Ilic, A.: Towards understanding the impact of personality traits on mobile app adoption - a scalable approach. In: Becker, J., vom Brocke, J., de Marco, M. (eds.) ECIS (2015), http://dblp.uni-trier.de/db/conf/ecis/ecis2015.html#XuFVI15