Using Statistical Information to Communicate Android Permission Risks to Users

Lydia Kraus, Ina Wechsung, Sebastian Möller
Quality and Usability Lab,
Telekom Innovation Laboratories, TU Berlin
Berlin, Germany
firstname.lastname@telekom.de

***Abstract*— The Android OS has a permission-based security system that controls the third party applications' access to sensitive information on the smartphone. The risk evaluation is left to the user who has to evaluate whether or not the requested permissions are appropriate. However, former work has shown that users lack attention to and understanding of the permissions which makes it difficult for them to make appropriate decisions. To support users with better understandable information we provide statistical information about permissions, grouped by functionality. We use methods from health risk communication to communicate this information to the users. In a lab experiment with 48 participants we find that users tend to choose more often the app with a lower number of permissions when statistical information is provided together with graphics. We also find that the privacy-intrusiveness and trustworthiness of apps is perceived differently when statistical information is given.**

***Keywords — Users; Android; permissions; app functionality; risk communication; statistical information*;**

I. INTRODUCTION

The Android platform is an operating system for mobile devices with fast growing popularity and a huge market for applications (apps). As of April 2013 there were about 850.000 apps available [1] in the official Android market, the Google Play Store. Android is open for third-party app providers and equipped with a permission-based security system to control app access to privacy- and security relevant resources. The permission-based system requires the user to evaluate if the required permissions are necessary.

Former work has shown that there are a number of usability problems with the current Android permission system - a large percentage of users are unaware of what permissions mean and what the possible consequences are [2]. Thus, most of the users might not be able to distinguish whether the requested permissions are indeed required to ensure the claimed functionality. Another drawback which makes it difficult for users to pay attention to permissions and thus to possible privacy issues is the time at which the permissions are shown to the user [3].

Currently the permissions are shown to the user when the decision for downloading an app has already been made, thus the user cannot include the number and quality of required permissions in the decision-making process. It has been shown that there are apps which ask for more permissions than would be necessary to ensure their functionality [4], [5]. However, excessive usage of permissions does not necessarily mean that an app is intentionally collecting data – it can also be a consequence of a programmer's lack of understanding on how to use the different permissions [4]. In some cases it might be even easier for developers to request plenty of permissions for their apps to make sure they work correctly [6]. Thus, the risk of permission abuse and data collection is difficult to determine for users and is subject to a high degree of uncertainty.

In this paper we introduce an approach to provide users with additional information in form of statistical data about the number of app permissions compared to other apps with similar functionality. The goal of our paper is to help users to easier interpret permission requests, to raise awareness of the permission issue, and to include the number of permissions in the decision-making process.

The paper is structured as follows: in Section 2, we provide an overview of related work regarding the user perspective, and the analysis of Android permissions and risk communication. In Section 3, we describe the research question. Section 4 presents the design and procedure of the user study while Section 5 provides an overview of the results. We conclude the paper with a discussion of the results, an overview of limitations, and future research in Section 6.

II. RELATED WORK

This section is divided into three parts. First, we give an overview of research conducted to communicate permission-related information to users. In the second part, we discuss former research about the categorization of Android apps with respect to permissions. As our approach is based on communicating risks to users by the help of statistical information, we conclude this section with previous work about communicating statistical information.

This is the accepted version of our STAST workshop paper published by IEEE under DOI 10.1109/STAST.2014.15.

A. *Communicating Android Permission Risks to Users*

Former work on users and Android permissions focuses on investigating user understanding of and attention to permissions [2], and on interface design to improve both [3, 7-9]. A fundamental paper in this field was published by Felt et al. in 2012 [2]. The work is based on the C-HIP model (Communication-Human Information Processing) which assumes the attention towards and the comprehension of warnings as prerequisites for motivation and behavior. In an online study they found that only a minority of users (17.5%) pay attention to permissions, and even less users understood all given permissions (2.6%). The authors conclude that the current permission system cannot be used in the way it is supposed to be used, as users can only make correct decisions if they understand what is being asked.

Other researchers focus on optimizing the representation of permissions and aim to include them in the decision-making process in order to facilitate the users' understanding of the permissions [3, 7]. The approach by Kelley et al. [3] is similar to ours: additional information about permissions is presented to users in the app market (before the decision is made) in the form of a privacy facts check list (a clustering of permissions in categories based on the kind of privacy sensitive information needed), and as a list of requested permissions. Thereby, the privacy check list leads to decreased installation rates of high-requesting apps, whereas only providing the list of permissions within the app store does not show significant effects on users' decision making behavior [3]. Hettig et al. [7] introduced another method for including permissions in the decision-making process by visualizing the accompanying risks with worst case examples. However, worst case examples are overestimated risks which might lead to a loss of user attention as consequently all apps are considered dangerous [2]. In contrast to our approach, permissions are not directly set into comparison to apps with similar functionality in the works of Kelley et al. [3] and Hettig et al. [7]. Benton et al. [8] found that providing an additional explanatory text about the permission use did not significantly influence users' installation behavior or regret of having installed a high-requesting app. On the contrary, adding visual cues to the provided information had a significant effect on users' installation behavior for some experimental conditions. Egelmann et al. [9] use a choice architecture to present apps of similar functionality and their requested permissions side-by-side, where low-requesting apps have a higher pricing than high-requesting apps. In a user study they find that when users are presented with a choice architecture, 25% of users are willing to pay a premium for privacy and put more weight on privacy as a decision factor. However, the space of a smartphone screen is limited, thus the number of apps with similar functionality which are presented side-by-side is limited as well; this issue could be solved by providing statistical information instead.

Another approach for risk notification is providing the user with additional information about permissions and privacy-intrusiveness of apps at install-time. Examples are described in Enck et al. [12] and implemented in Clueful Privacy Advisor [10] and Androlyzer [11]. There are also apps which help to protect the users' privacy in real-time by policy enforcement [13]. For enforcing (privacy) policies, it has to be ensured that the "treated" app is still working correctly even if certain permissions are switched off manually. Otherwise the user has to sacrifice functionality for security. Another method to make privacy-intrusiveness more transparent for users is monitoring app behavior in real-time [14]. However, we will not take into account approaches where the privacy-intrusiveness of an app is determined after installing an app or at the time of the installation, as we focus on informing the users before a decision is made. Once a user is using an app, the privacy might have been already intruded. Also other effects, such as social interaction or good usability might appear more valuable to the users than privacy-intrusiveness.

B. *Analyzing Android permissions*

We overviewed several works on the automatic evaluation of permission requests [4, 6, 15-17] in order to find statistics upon which we can build user interfaces.

Felt et al. [4] analyzed automatically the code of 940 Android applications and found that one third of the tested apps requested more permissions than necessary. However, they provide no statistics about permission usage by app functionality. In two papers relying on machine learning techniques to identify potential malicious [6] or low-quality apps [17] statistics about permissions are given; these are, e.g. the average number of permissions by Google Play category or the Android permissions requested most frequently. The mentioned approaches [6, 17] are based on the predefined Google Play categories. The functionalities of apps within these categories can strongly vary which makes statements about unusual permissions for an app of a certain category difficult. Furthermore, approaches for automatic classification of Android apps based on the requested permissions can be found in the literature [15, 16]. However, these are inverse to our approach as they infer functionality from permission patterns.

In summary, none of the above mentioned approaches suited our purpose well enough. Therefore, we decided to manually extract permission statistics for three functionalities (cf. Section 4).

C. *Communicating statistical information*

One of the common techniques to visualize descriptive statistics is the box plot [18]. Also in the health domain statistical information about risks (e.g. risk to develop cancer) has to be communicated to lay users. In this context it was shown that communicating risks in "natural frequencies" (i.e. using integers to communicate frequencies instead of percentages) supports risk understanding and calculation with risks [19]. Visual communications of risks is an additional way to support risk understanding and evaluation. Graphics help in revealing data patterns, in supporting mathematical operations (e.g. comparison), and in attracting and holding people's attention [20]. Lipkus and Hollands [20] give a detailed overview of ways to communicate risks visually, including risk ladders, line graphs, bar graphs, pie charts, and histograms.

III. Communicating Android permission risks by providing statistical information in the app market

A. Communicating statistical information about permissions

When browsing Google Play Store we noticed that there are huge differences in permissions between apps with same or similar functionalities. Thus, based on results from related work, we will examine if statistical information could help users to compare apps with similar functionality and include privacy-intrusiveness as a decision criterion. We assume that presenting statistics is a suitable method to communicate the uncertainties that are associated with the permission system. Moreover, statistics give users the freedom to evaluate the risk on their own without pushing them in one or the other direction by defining which threshold is "good" or "bad". Also, we assume that with our approach it is possible to support users' decision-making without requiring them to understand the exact meaning and implications of each permission.

We decided to collect statistical information on the number of permissions (minimum, maximum, mean, median, 1st and 3rd quartile) for three types of apps (functionalities): weather forecasts (weather), torches (torch), and memory games (memory). The reason for this selection was the high number of apps providing these functionalities which helps to provide a rich statistic. We also assumed that the functionalities are easy to understand by users, as most of the people should know what a weather forecast, a torch and a memory game are. We are aware that the pure number of permissions is still not enough to exactly evaluate the privacy-intrusiveness of an app and that in future studies also the permission type should be taken into account. The presented interfaces are a first approach to determine how statistical information in the context of Android permissions affects users and their decision making.

B. User Interfaces

1) Permission statistics by app functionality.

To collect statistical data about permissions by app functionality we accessed the German version of the Google Play Store in July 2013 with a Sony Xperia S smartphone. To find weather forecast apps we entered the keyword "weather forecasts" in the app search, for torch apps "torch" and for memory games "memory". We decided to scan the first 200 results of each search. We only selected free apps and included only those with basic functionalities e.g. for weather apps features like forecast, rain radar and widgets were selected, but apps with special features such as wind, boating, and bike forecasts, as well as weather alerts were excluded.

For torch we selected apps with all kind of lightning functionalities and also flashlights, and for memory games we selected apps with and without online score functionality. After removing apps with non-basic functionalities from the initial list (N=200 for each category) we were left with 111 apps for weather, 192 for torch, and 133 for memory (for summary statistics cf. Table 1).

We collected different ideas for the interfaces in an ideation session. It was important for us to see how people react to pure textual information and to information that is a mix of text and graphic. It would have been an option to present all the information which we collected in a boxplot; however a boxplot might be difficult to understand for lay users. Also we decided to exclude textual information about median and 3rd quartile as we wanted to avoid overloading users with information.

TABLE 1. Descriptive statistics of number of permissions for the three app categories

Category	**Permission statistics**					
	Min.	***1st Qu.***	***Median***	***Mean***	***3rd Qu.***	***Max.***
Weather	1	3	5	5.53	7	18
Torch	0	3	5	5.14	6	19
Memory	0	3	4	5.84	7	16

2) Standard UI

We tested three different user interfaces. The first called "Standard UI" (cf. Fig. 1.) is of similar design as the Google Play Store, and it was our control condition.

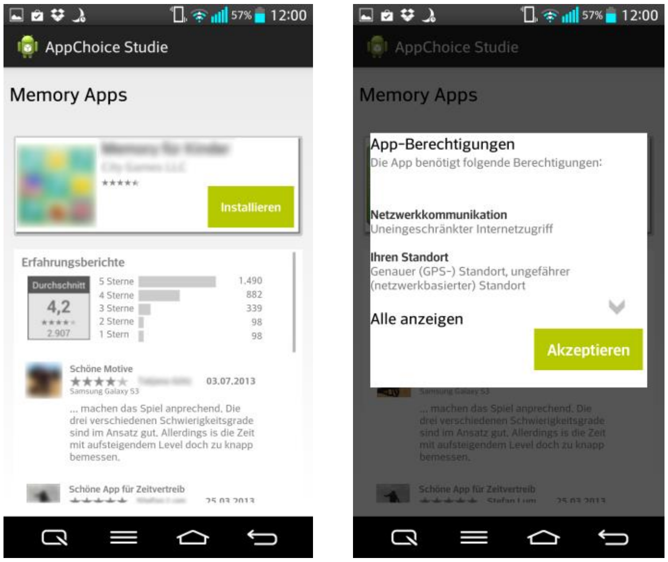

... machen das Spiel anprechend. Die drei verschiedenen Schwierigkeitsgrade sind im Ansatz gut. Allerdings is die Zeit mit aufsteigendem Level doch zu knapp bemessen.

App-Berechtigungen
Die App benötigt folgende Berechtigungen:

Netzwerkkommunikation
Uneingeschränkter Internetzugriff

Ihren Standort
Genauer (GPS-) Standort, ungefährer (netzwerkbasierter) Standort

Fig. 1. Original "Standard UI" (left). After pressing the installation button, the permission dialog appears (right).

The "Standard UI" provided the user with the most important information including three screenshots of an app, a textual description of the app, star ratings, number of downloads, and six reviews of other users. The permissions of an app are only visible in a pop-up window after the "install" button is pressed, as it is current practice in Google Play.

The second user interface called "Text UI" (cf. Fig. 2.) and the third called "Graphic UI" (cf. Fig. 2.) provide additional statistical information beneath the app description, thus information about permissions is shown before pushing the "install" button.

3) Text UI

In the textual prototype we provided the users with the number of permissions of the current app, the mean, and the 1st quartile of the permission statistics of the category of the current app. Based on Hoffrage et al. [19] who stated that it is helpful when statistics are communicated to people in natural numbers, we decided to provide the 1st quartile information in natural numbers (using the term “25 of 100 apps” instead of “25% of apps”). We placed the information directly below the app description. Also, we put the list of permissions of the current app below the textual information. After pressing the “install” button, users see again the pop-up window with the permissions which they need to accept in order to install the app, as in the “Standard UI”.

4) Graphic UI

In the graphical prototype we provided the user with text and a graphic about the number of permissions of the current app, the minimum, the maximum, the mean, and the 1st quartile (cf. Fig. 2).

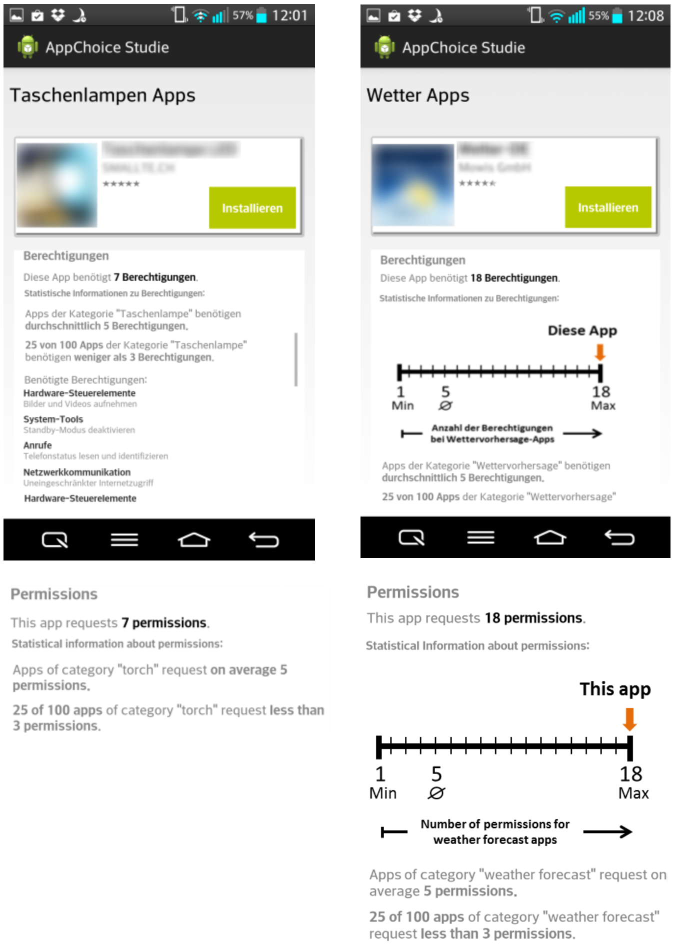

Taschenlampen Apps

Berechtigungen

Diese App benötigt **7 Berechtigungen**.

Statistische Informationen zu Berechtigungen:

Apps der Kategorie "Taschenlampe" benötigen **durchschnittlich 5 Berechtigungen**.

25 von 100 Apps der Kategorie "Taschenlampe" benötigen **weniger als 3 Berechtigungen**.

Benötigte Berechtigungen:

Hardware-Steuerelemente
Bilder und Videos aufnehmen

System-Tools
Standby-Modus deaktivieren

Anrufe
Telefonstatus lesen und identifizieren

Netzwerkkommunikation
Uneingeschränkter Internetzugriff

Hardware-Steuerelemente

Wetter Apps

Berechtigungen

Diese App benötigt **18 Berechtigungen**.

Statistische Informationen zu Berechtigungen:

Apps der Kategorie "Wettervorhersage" benötigen **durchschnittlich 5 Berechtigungen**.

25 von 100 Apps der Kategorie "Wettervorhersage"

Permissions

This app requests **7 permissions**.

Statistical information about permissions:

Apps of category "torch" request **on average 5 permissions**.

25 of 100 apps of category "torch" request **less than 3 permissions**.

Permissions

This app requests **18 permissions**.

Statistical Information about permissions:

Apps of category "weather forecast" request on average **5 permissions**.

25 of 100 apps of category "weather forecast" request **less than 3 permissions**.

Fig. 2. Original “Text UI” (upper left), original “Graphic UI” (upper right), translation of “Text UI” (lower left), and translation of “Graphic UI” (lower right).

We decided to present visual information in the “Graphic UI” in the form of a horizontal risk ladder. We used a horizontal shape as it allowed us to show additional textual information without the necessity of scrolling. Former research has shown that risk ladders help people anchoring a risk more effectively as it provides them with upper and lower reference points [20]; moreover, lengths are usually perceived by people without bias [20]. We decided to use a linear scale from the minimum to the maximum number of permissions used in one category. The text below the risk ladder was the same as in the “Text UI”. Also the list of permissions was given below the text.

C. Hypotheses

To assess users’ perception regarding the developed interfaces, we conducted a user study. The study was designed to test the below-mentioned hypotheses. Our main interest is to determine if the users’ willingness to download an app is influenced by additional statistical information about the number of app permissions. We know from related work that graphics can be more powerful than text in communicating risks [20], hence we expect that while the “Text UI” will influence the users, this influence will not to be as strong as for the “Graphic UI”. Thus, we assume that the magnitude of the effects will be ordered “Standard UI “< “Text UI” < “Graphic UI”:

H_{1a}: Participants will decide more often for the low-permission app when presented with “Graphic UI” compared to “Standard UI”.

H_{1b}: Participants will decide more often for the low-permission app when presented with “Text UI” compared to “Standard UI”.

H_{1c}: Participants will decide more often for the low-permission app when presented with “Graphic UI” compared to “Text UI”.

We also expect that the permission ratio between the high-permission and the low-permission app influences the app installation rate. Within the experiment we selected a high ratio for the weather category, a medium ratio for the torch category and a small ratio for the memory category:

H_{2a}: For the high ratio category (weather) participants will more often install the low-permission app in all UIs compared to the low and medium ratios (due to the extreme difference in the number of permissions).

H_{2b}: For the medium ratio category (torch) participants will more often install the low-permission app compared to the low ratio category.

Furthermore we want to determine if the “Text UI” and the “Graphic UI” influence users regarding the importance of the number of permissions as a criterion for the app selection. Again, we assume that the magnitude of the effect will be ordered “Standard UI” < “Text UI” < “Graphic UI”:

H_{3a}: When presented with the “Graphic UI”, participants will rate the importance of app permissions as a decision factor higher compared to the “Standard UI”.

H_{3b}: When presented with the “Text UI”, participants will rate the importance of app permissions as a decision factor higher compared to the “Standard UI”.

H_{3c}: When presented with the “Graphic UI”, participants will rate the importance of app permissions as a decision factor higher compared to the “Text UI”.

To see if our approach is able to raise awareness regarding possible permission abuse, we asked participants to rate the perceived privacy of and trust in an app on a continuous scale.

Thus, we can find out whether participants connect the number of permissions with the privacy-intrusiveness of an app and trust in an app.

H_{4a}: The low-permission app will receive higher perceived privacy ratings compared to the high-permission app.

H_{4b}: The low-permission app will receive higher trust ratings compared to the high-permission app.

H_{4c}: There will be an interaction effect between the user interfaces and the perceived privacy and trust. Regarding the high-permission app we expect lower privacy and trust ratings for the "Text UI" and the "Graphic UI" and vice versa for the low-permission app compared to the "Standard UI".

To investigate whether people with higher privacy concern are more sensitive regarding the "Text UI" and the "Graphic UI" we also used the Global Information Privacy Concern scale (cf. procedure section).

IV. EXPERIMENTAL SETUP

The study was implemented as a lab experiment in which we required participants to bring their own smartphone in order to create a more realistic setting. We chose a within-subjects design to increase the power of the statistical tests. The experiment consisted of three parts so each user had to make three decisions. In each part we showed one of the UIs and presented them with two different apps of the same functionality, one app with a high number of permissions and the other one with a lower number of permissions. After both apps were presented, we asked the participants to decide for one app and to install it on their phone

An app called "AppChoice" (designed for the experiment) led the participants through the experiment and presented them with our custom app store. In order to counterbalance possible influencing factors (e.g. user reviews, graphical design, user ratings, number of downloads), we decided to use mock-up apps instead of real apps, but the participants were unaware of this fact. To maintain the impression that the participants were confronted with real apps, we simulated a push-message (a moving arrow-icon shown when apps are downloaded together with the message "app is being installed") after the installation button was pressed. To avoid that participants uncover our story we asked them not to navigate away from AppChoice during the experiment. To evaluate whether there is an effect of the permission ratio (the ratio between the number of low permissions and the number of high permissions within the functionalities) we selected different permission ratios for each functionality. We randomly selected weather to have the high ratio, torch to have the medium ratio and memory to have the low ratio, as the descriptive statistics of all three functionalities looked similar (cf. Table 1). Therefore, we chose as a high ratio for weather 18 (maximum) vs. 4 (interquartile range) permissions, as a medium ratio for torch 7 (interquartile) vs. 2 (below 1st quartile) and as a low ratio for memory 4 (interquartile) vs. 0 (minimum) permissions.

The mock-up apps we used were inspired by real apps. We took pictures from low ranked apps of Google Play (less than rank 50 on the search results), to avoid asking people to download apps which they already might have installed. We selected two sets (a set of low permissions and a set of high permissions) of typical permissions for each functionality (app type) and two sets of user reviews from highly ranked apps for the related functionality (for the distribution of star ratings of the reviews cf. Table 2).

The order of appearance of the two apps with the same functionality on the overview page was randomized. We also counterbalanced the order of the functionalities (weather, torch, memory). The user interfaces were always shown in the order "Standard UI" – "Text UI" – "Graphic UI". As the salience of the permissions increases in strength from "Standard UI" to "Graphic UI" we always chose this order to avoid priming effects.

TABLE 2. USER REVIEW RATING DISTRIBUTION AMONG APP CATEGORIES

Category	**User review ratings**				
	5-star	***4-star***	***3-star***	***2-star***	***1-star***
Weather	2	1	2	-	1
Torch	3	1	1	-	1
Memory	2	3	1	-	-

A. *Procedure*

Participants were first given a demographic questionnaire. Then they were asked to rate the importance of eight factors for their decision on a 7 point scale (1 = not important at all, 7 = very important). The decision factors were: description of the app, visual impression of the app, reviews provided by other users, ratings (number of stars), number of downloads, permissions requested by the app, provided functionality (according to description), publisher (company). After each installation they were again asked to fill in the questionnaire about the decision factors. To measure perceived privacy of and the trust in both, the selected and the not selected app, a continuous scale ranging from "low" over "medium" to "high" (min. = 0, max. = 21) was used. To cover up that privacy and trust were the most important items for us we mixed in four questions regarding the overall rating of the apps and impressions about aesthetics and the like. After the 3rd and last installation they were presented with a questionnaire about privacy concerns. We took the Global Information Privacy Concern scale provided in Malhotra et al. [21]. We logged all interactions with AppChoice during the experiment. Participants were asked for consent before the experiment, and were debriefed at the end of the experiment. When the participants arrived in our lab we told them that the objective of the experiment is to investigate people's impressions of Android apps and we did not mention any privacy or security related issues. After we installed AppChoice on their device we told them that the presented apps are from the "app store". Within the experiment the participants were supposed to decide for one of the apps with the same functionality after an exploration phase of maximum 5 minutes. We also informed them that they may uninstall all apps that were installed during the study after the study has finished.

B. *Participants*

Participants for our study were recruited with classified online ads on Ebay, classified ads of a city magazine and a subjects-portal of our university. We required German-speaking Android phone owners willing to install Android apps on their own device during the study. They were paid 15€ for 60-90 minutes. Our sample was 50% female. Participants were between 18 and 60 years old with an average of 31.68 years (SD = 11.70). All kind of occupation groups were covered including 14 employees (29.2%), four self-employed (8.3%), 15 students (31.3%), three apprentices (6.2%), four pupils (8.3 %), two pensioners (4.2%), one housewife or stay at home husband (2.1%), three unemployed (6.2%) and two others (4.2 %). 16 (33.3%) of our participants had less than a high school degree, 19 (39.6 %) had a high school degree and 13 (27.1%) had a university degree. All participants were Android users and used phones with Android versions 2.1, 2.2, 2.3, 4.0, 4.1 and 4.2.

V. Results

In this section we report the results of the experiment. We analyze the quantitative data we obtained regarding the installation rates, the importance of the number of permissions as a decision factor, the perceived privacy of and trust in the low- and high-permission app, as well as the importance of the other decision factors.

A. *Installation rates*

1) *Comparing the three decisions of all participants.*

We used a Cochran's Q test to compare the number of low- and high-permission installations (hereafter referred to as “installation rate”) for the three decisions of all 48 participants. We found a significant difference in the installation rate of the low-permission app between the three types of user interfaces, $Q(\text{df} = 2, N=48) = 6.07$, $p_{1\text{-tailed}} = 0.03$. A pairwise comparison using McNemar's test with Bonferroni- correction revealed that significantly more participants selected the low-permission app in the “Graphic UI” than in the “Standard UI” (cf. Table 3). The differences between “Text UI” and “Standard UI” and between “Graphic UI” and “Text UI” were not significant.

TABLE 3. Installation rates for the low-permission app.

	Pairwise comparison of installation rates						
	Stand. UI	***Text UI***	***Graph. UI***	***Diff. Text/ Stand.***	***p-Val. Text / Stand.***	***Diff. Graphic/ Stand.***	***p-Val. Graphic/ Stand.***
n	27	32	38	+5	0.202	**+11**	**0.01**
%	56.25	66.67	79.17	+10.42		**+22.92**[a]	

[a] Significant differences in bold, Bonferroni-corrected. *n* is the number of participants who chose the low-permission app.

2) *Installation rate as a function of the permission ratio.*

We used Fisher’s exact test to determine whether or not there are differences in the installation rate between the three permission ratios (operationalized as weather, torch and memory). Thus, for each UI, the installation rate of the 16 participants who were presented with the weather app was compared with the installation rate of the 16 participants who were presented with the torch app and the installation rate of the 16 participants who were presented with the memory app. Although we expected the high and the medium permission ratio to have an influence on the installation rate, we did not find a significant difference.

3) *Installation rate as a function of privacy concern.*

For all UIs, a χ^2-test showed no difference in the app installation rate between participants for low, medium, and high privacy concern.

B. *Importance of the number of permissions as a factor for the app selection*

We used a repeated-measure ANOVA to compare the mean importance of the number of permissions as a decision factor for the three decisions of all 48 participants. We found that the importance of the number of permissions as a decision factor differed significantly between the UIs, $F(2, 90) = 22.01$, $p<0.01$, part. $\eta^2 = .328$. A Sidak-corrected post-hoc analysis showed that the importance in the “Standard UI” was significantly lower than the importance in the “Text UI” and the “Graphic UI”. Between “Text UI” and “Graphic UI” the post-hoc test indicated no significant difference.

TABLE 4. Mean values of importance of number of permissions as decision criterion.

Pairwise comparison of importance of permissions						
Stand. UI	***Text UI***	***Graphic UI***	***Diff. Text/ Stand.***	***p-Val. Text / Stand.***	***Diff. Graphic/ Stand.***	***p-Val. Graphic/ Stand.***
3.41 (2.21)	5.11 (2.10)	5.41 (2.04)	**+1.7** [b]	**<0.01**	**+2.0** [b]	**<0.01**

[b] rated on a 7-point scale, where 7 is “very important”. Significant differences in bold, Sidak-corrected. Standard deviations in brackets.

1) *Interrelation between number of permissions, perceived privacy and trust*

We used a repeated-measure ANOVA, with UI and permission level (low and high) as within-factors, to compare the difference in perceived privacy between the low-permission app and the high-permission app for each UI. We found a significant difference for perceived privacy between the low-permission app and the high-permission app, $F\ (1, 36) = 55.97$, $p<0.01$, *part.* $\eta^2 = .609$. A pairwise comparison using Bonferroni-correction revealed that for the “Text UI” and the “Graphic UI” the low-permission app had a significantly higher perceived privacy than the high-permission app (cf. Fig 3 and Table 5).

We also found an interaction effect between UI and permission level, $F\ (2, 72) = 14.21$, $p<0.01$, part. $\eta^2 = .283$. A pairwise comparison using Bonferroni-correction revealed that the low-permission app had a significantly higher perceived privacy for the “Text UI” (p=0.011) and the “Graphic UI” ($p<0.01$) compared to the “Standard UI” (cf. Fig. 3). The high-permission app had a significantly lower perceived privacy for the “Text UI” ($p<0.01$) compared to the “Standard UI” (cf. Fig. 3).

To compare the difference of trust in the low-permission app and the high-permission app for each UI we used again a repeated-measure ANOVA with the UI and permission level as within-factors. We found a significant difference between trust in the low-permission app and the high-permission app, $F(1, 40) = 36.05$, $p<0.01$, *part.* $\eta^2 = .474$. A pairwise comparison using Bonferroni-correction revealed that for the “Text UI” and the “Graphic UI” the low-permission app had a significantly higher trust rating than the high-permission app (cf. Fig. 3. and Table 5). We also found an interaction effect between UI and permission level, $F(2, 80) = 5.56$, $p<0.01$, *part.* $\eta^2 = .122$. A pairwise comparison with Bonferroni-correction revealed that the low-permission app had a significantly higher trust rating for the “Graphic UI” ($p<0.01$) compared to the “Standard UI” (cf. Fig. 3).

TABLE 5. MEAN VALUES OF PERCEIVED PRIVACY AND TRUST.

Pairwise comparison of perceived privacy				
UI	***Perm.***	***Perceived Privacy/ Trust***	***Diff. high-low***	***p-Val.***
Stand. UI	low	μ =14.92/ 14.71 (σ = 5.30/ 5.53)	Priv.:-1.42 Trust: -0.59	Priv.: 0.086 Trust: 0.64
	high	μ = 13.5/ 14.12 (σ = 4.79/ 5.41)		
Text UI	low	μ = 17.67/ 17.30 (σ = 5.13/ 5.25)	**Priv.:-9.24** [c] **Trust:-7.06** [c]	**Priv.:<0.01** **Trust:<0.01**
	high	μ =8.43/ 10.24 (σ = 5.91/ 5.92)		
Graph. UI	low	μ = 17.74/ 17.59 (σ = 4.82/ 4.44)	**Priv.:-6.87** [c] **Trust:-6.57** [c]	**Priv.:<0.01** **Trust:<0.01**
	high	μ = 10.87/ 11.02 (σ = 6.49/ 5.62)		

[c] rated on a continuous scale from 0 to 21 (0 = low; 21 = high). Significant differences between means in bold, pairwise comparison, Bonferroni-corrected.

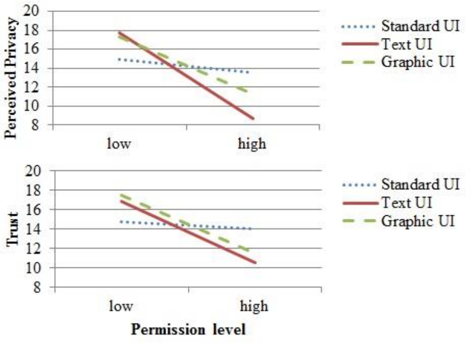

Fig. 3. Comparison of perceived privacy and trust for the low-permission app and the high-permission app by user interface

C. Importance of other decision factors

In addition to the importance of the number of permissions, other decision factors were analyzed. Rated on a 7-point scale from 1 (not important at all) to 7 (very important), description, functionality, user reviews, visual impression and ratings were rather important for the participants and received mean values between 5.09 and 5.69. Publisher (2.04-2.15) and number of downloads (3.26-3.48) were not as important for the participants. A pairwise comparison of the decision factors with Bonferroni-correction showed no significant difference for all factors (except permissions, cf. paragraph 5.2) between the UIs.

VI. DISCUSSION

The results of our work suggest that statistical information can affect users in their decision making and perception of privacy when selecting Android apps. When statistical information was provided in form of the “Graphic UI”, participants decided significantly more often to choose the low-permission app. Thus, H_{1a} was supported, but H_{1b} and H_{1c} were not supported. This indicates that if statistical information is included in the app market, it needs to be presented in an attention-catching way in order to influence users’ behavior. The results are in line with [3] and [8]. Kelley et al. [3] found that including a list of permissions in the app description before the download-decision is made did not show significant effects on users’ downloading behavior. This is similar to our “Text UI” condition, even though we had more text (the statistical information) to explain the permission use. Also, Benton et al. [8] found that providing additional explanation text about permission use did not significantly influence the installation rate, but adding visual cues to the provided information had a significant effect on users’ installation behavior for some experimental conditions.

Contrary to our expectations, the permission ratio did not show significant effects on the installation rate: therefore, H_{2a} and H_{2b} were not supported.

Overall, in the “Text UI” condition and the “Graphic UI” condition the additional statistical information influenced users to increase the importance of the number of permissions as a decision factor. Thus, H_{3a} and H_{3b} were supported, but H_{3c} was not supported. This result indicates that providing comparative information before the decision is made helps the users to include the number of permissions in their decision making. This is in line with results of Egelmann et al. [9] who found that providing users with comparative information during the decision making led to higher valuation of privacy.

Statistical information in both, the “Text UI” and the “Graphic UI”, also influenced users’ perceived privacy of and trust in an app with respect to the number of permissions. With statistical information given, participants perceived apps with a higher number of permissions less privacy-protecting and less trustworthy compared to the apps with a lower number of permissions. Thus, both H_{4a} and H_{4b} were supported for the “Text UI” and the “Graphic UI”. H_{4c} was supported as well. This suggests that additional statistical information can help to raise awareness with respect to permission abuse.

No significant differences for the installation rate were found between participants with low, medium and high privacy concern which suggests that participants were sensitive to the information, irrespective of their general concern. This is in line with results of Egelmann et al. [9], who also did not find a significant difference in the installation behavior of participants with low, medium and high privacy concern.

Other decision factors (cf. paragraph C in the results section) were not significantly influenced between the decisions. This suggests that providing additional comparative information about app permissions does not lead participants to exclude other decision criteria.

A. Limitations and Future Work

The experiment was intended to get a first impression on how statistical information in the app market affects users' decision making and perceived privacy of apps. Implementing the study as a lab experiment allowed us to have a controlled setting and to give participants the impression of actually being in a choice situation with real risk, but it also limited us to a smaller sample size compared to online studies. Due to the small sample size we decided for a within-subjects design to increase the power of the statistical tests. We note that presenting the user interfaces in the same order for all participants might have led to learning effects. However, randomizing the order would have possibly led to priming effects as the salience of the permissions was higher for the "Graphic UI" compared to both other UIs. Nevertheless, the constant order of the UIs limits the validity of the results; in future studies a between-design could be chosen to circumvent this drawback.

For the interface design we only considered the number of permissions and not the kind of the permissions themselves, i.e. whether some permissions cause stronger effects than others. Including this information should be a subject for future studies.

There are some limitations on the collection of statistical data in the app market in general. For this study we collected manually statistical information about apps with similar functionalities. In order to apply the concept broadly, an automated approach needs to be developed. Also, when apps with completely new functionalities enter the market, it is difficult or impossible to collect statistical information. The same is true for apps of similar functionality with only few samples in the market.

Our approach might rather support users who search for apps with specific functionalities than those who search for a specific app. Moreover, letting participants choose between only two apps and not between several apps is a limitation and should be addressed in future studies.

B. Conclusion

In this paper we introduced an approach to provide users with additional information in form of statistical data about the number of app permissions in relation to other apps with similar functionality. The goal was to help users to easier interpret permission requests, to raise awareness regarding the permission issue, and to include the number of permissions in the decision-making process. The results of a user study suggest that when statistical information is provided, users tend to choose more often apps with a lower number of permissions; that they perceive low-permission apps as less privacy-intrusive and more trustworthy; and that they actively started to include the number of permissions in the decision-making process.

ACKNOWLEDGMENT

This work has been supported by the EU FP-7 support action ATTPS under grant agreement no. 317665. We would like to thank Robert Schmidt for implementing the AppChoice app and helping to conduct the study, Sheikh M. Habib for his helpful comments, and Maija Poikela and Jens Ahrens for proofreading the paper. Also, we would like to thank all anonymous reviewers for their helpful comments.

REFERENCES

[1] Statista – Das Statistik Portal, http://de.statista.com/statistik/daten/studie/74368/umfrage/anzahl-der-verfuegbaren-apps-im-google-play-store/ (accessed April 29, 2014)

[2] Porter Felt, A. ; Ha, E.; Egelmann, S.; Haney, A.; Chin, E. & Wagner, D.: Android Permissions: User Attention, Comprehension, and Behavior, *Symposium on Usable Privacy and Security, Washington, DC, USA*, 2012

[3] Kelley, P. G.; Cranor, L. F. & Sadeh, N.: Privacy as Part of the App Decision-Making Process, *CHI 2013*, 2013

[4] Porter Felt, A. ; Chin, E.; Hanna, S.; Song, D. & Wagner, D.: Android Permissions Demystified, *Proceedings of the ACM CCS, Chicago, Illinois, USA*, 2011

[5] The New York Times, http://www.nytimes.com/2012/10/29/technology/mobile-apps-have-a-ravenous-ability-to-collect-personal-data.html?_r=1& (accessed April 29, 2014)

[6] Barrera, D.; Kayacik, H. G.; Van Oorschot, P.C. & Somayaji, A.: A Methodology for Empirical Analysis of Permission-Based Security Models and its Application to Android, *Proceedings of the ACM CCS, Chicago, Illinois, USA,* 2010

[7] Hettig, M.; Kiss, E.; Kassel, J.-F.; Weber, S.; Harbach, M. & Smith, M.: Visualizing Risk by Example: Demonstrating Threats Arising From Android Apps, *Symposium on Usable Privacy and Security (SOUPS), Newcastle, UK,* 2013

[8] Benton, K., Camp, L. J., & Garg, V. (2013, March). Studying the effectiveness of android application permissions requests. *IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops),* 2013

[9] Egelman, S., Felt, A. P., & Wagner, D.: Choice architecture and smartphone privacy: There's a price for that. In *The Economics of Information Security and Privacy* (pp. 211-236). Springer Berlin Heidelberg, 2013

[10] Clueful Privacy Advisor on Google Play, https://play.google.com/store/apps/details?id=com.bitdefender.clueful&hl=de (accessed April 29, 2014)

[11] Androlyzer –Know more about your apps, https://www.androlyzer.com/ (accessed April 29, 2014)

[12] Enck, W.; Ongtang, M. & McDaniel, P.: On Leightweight Mobile Phone Application Certification, *Proceedings of the ACM CCS, Chicago, Illinois, USA*, 2009

[13] Backes, M.; Gerling, S.; Hammer, C.; Maffei, M. & von Styp-Rekowsky, P.: AppGuard - Real-time policy enforcement of third-party applications, Saarland University, 2012

[14] Bal, G.: Revealing Privacy-Impacting Behavior Patterns of Smartphone Applications, *Workshop on Mobile Security Technologies (MoST), San Francisco, USA,* 2012

[15] Sanz, B.; Santos, I.; Laorden, C.; Ugarte-Pedrero, X. & Bringas, P. G.: On the automatic categorisation of android applications, *Consumer Communications and Networking Conference (CCNC),* 2012

[16] Rassameeroj, I. & Tanahashi, Y.: Various Approaches in Analyzing Android Applications with its Permission-Based Security Models, *IEEE International Conference on Electro/Information Technology,* 2011

[17] Frank, M.; Dong, B.; Felt, A. P. & Song, D.: Mining Permission Request Patterns from Android and Facebook Applications, *IEEE 12th International Conference on Data Mining (ICDM),* 2012

[18] Potter, K.: Methods for presenting statistical information: The box plot. Visualization of Large and Unstructured Data Sets, (LNI) 4, 2006

[19] Hoffrage, U.; Lindsey, S.; Hertwig, R. & Gigerenzer, G.: Communicating Statistical Information, Science, 2000, 290, 2261-2262

[20] Lipkus I.M., Hollands J.G.: The Visual Communication of Risk, Journal of the National Cancer Institute Monographs No. 25, 1999

[21] Malhotra, N. K.; Kim, S. S. & Agarwal, J.: Internet Users' Information Privacy Concern (IUIPC): The Construct, the Scale and a Causal Model. Information Systems Research, 2004, 15, 336-355