Can You Hear Me Now?: A Technical Report on Combining Audio with Privacy Permissions

SHAKTHIDHAR REDDY GOPAVARAM, Indiana University
OMKAR BHIDE, Indiana University
JEAN CAMP, Indiana University

The Android and iOS privacy ecosystems are grounded in the permissions, which provide information to and control of access to phone resources. These are instrumented with a combination of permission manifests at the time an app is selected, resource warnings at first use, and per-resource controls. As yet the controls provided in the form of permissions has proven insufficient to address privacy concerns and prevent selection of malicious apps. Here we alter the visual presentation, the media of the presentation, and the timing of permissions information. The goal is to make it simple to compare permissions (and thus privacy) of different apps when the person is focused on the task of comparing and selecting apps. We provide aggregate ratings which allow for simple comparisons, and add a small audio feedback component. Our goal is timely, comprehensible, cognitively simple permissions.

Specifically we test risk information with padlock icons and short audio notifications. The combined sound and icons were significant. Overall, these simple privacy cues has a consistent effect on app choices. Adding sound as a form of feedback and simplifying comparison of apps makes significant change in individual decision-making. In the aggregate even small biases towards more privacy in individual apps and in individual phones could have an impact on the Android ecosystem.

CCS Concepts: •**Security and privacy** → **Usability in security and privacy;**

Additional Key Words and Phrases: Privacy; Usable security; Human Factors; Smartphones; Android; Simulation

ACM Reference Format:
Shakthidhar Reddy Gopavaram, Omkar Bhide, Jean Camp, 2017. Can You Hear Me Now?: A Technical Report on Combining Audio with Privacy Permissions *ACM Trans. Embedd. Comput. Syst.* 9, 4, Article 39 (March 2010), 20 pages.
DOI: 0000001.0000001

1. INTRODUCTION

Smartphone apps potentially store and have access to sensitive data (like personal information, location information, contacts, financial information etc.) and sensor data (like cameras, microphones, accelerometers etc.). Information exfiltration, both as criminal attacks and through careless leakage, is severe a threat for mobile users and the organizations where they work, bank, or otherwise engage. The risk associated with a given resource or permission is often unclear to the individual or the organization. In fact, researchers (and criminals) are sufficiently innovative that the practical risk of permissions can be difficult to identify. For example, researchers have demonstrated how a malicious application could use an accelerometer to decode users keystrokes [Marquardt et al. 2011; Owusu et al. 2012], which could be used to learn users password and other sensitive information.

OS providers of popular mobile platforms understand the severity of the threat posed by cybercriminals to smartphone devices and are taking steps to mitigate risk. (In this paper "OS providers" refers to Apple for iOS, Google for Android, Microsoft or Windows, and Amazon for Fire). Google, for example, initially allowed app developers to publish their apps on the Play Store without initial review. This is no longer the case; every app that is submitted to the Play Store is tested and reviewed to identify policy violations and malware.

Despite the efforts of OS providers, malicious applications are still finding their way onto app stores. Recently, the Indian government launched an Android app, BHIM (Bharath Interface for Money), to promote cashless transactions in India. But, according to news reports, there are over three dozen fake BHIM apps that have popped up on Googles Play store [Bali 2017]. These fake apps are compromising the personal and financial security of the people who unsuspectingly downloaded them. These kinds of masquerading attacks have occurred on iOS too (A lot of fake retail apps popped up on Apples App Store before the holidays) [Goel 2016].

Another problem faced in todays smartphone ecosystem is the usage of personal information in direct opposition to people's expressed preferences. Prior research has indicated that smartphone users are mostly unaware of the data collected by the apps installed on their phones and when they learn about what data is being collected and how it is being used they express severe concerns [Acquisti et al. 2015a]. In one case, participants felt deceived and expressed severe concerns when they learnt about data collected by the fruit ninja app [Shklovski et al. 2014]. Access to personal information by apps has been found in both iOS [Egele et al. 2011; Agarwal and Hall 2013] and Android [Enck et al. 2011; Beresford et al. 2011; Enck et al. 2014; Zhou et al. 2011; Arzt et al. 2014].

Mobile privacy impinges the physical world in a way that desktop privacy does not. Consider, for example, location in the case of two apps enabled by the weaknesses in privacy controls, the first of which was blocked by Foursquare. *Girls Around Me* displayed the Facebook profiles of any women, including any public details on Facebook and tagged information such as pictures, of women in at a specified location. The app advertised *Creepy (Location Creeper)* is self-explanatory, it tracks either identified individuals as they move around or it tracks individuals moving around in adjacent space. Both tracking apps create risks, particularly since tracking apps are an issue for those at risk for relationship violence or stalking. In the physical domain a lack of privacy can correlate with actual physical risk. The targeted people did not download the app, and the risk of their information sharing was invisible. Google tried to address this by identifying particular permissions as more of less sensitive; however, this has not been found to be highly useful in individual decision-making.

In order to alleviate these problems, OS providers leverage privacy frameworks and standards to ensure that users information is aggregated and stored in a privacy preserving manner [Sweeney 2002]. Yet different apps use different business models and often request permissions promiscuously for ease of development or interoperability with different advertising networks.

Since prevention of privacy violations committed either by legitimate or by malicious applications is not entirely plausible, the next best option we have is to communicate this information to the end user so that he or she can make informed decisions. Currently, the two leading mobile operating systems (iOS and Android) use warning prompts or install time permissions manifests to convey the privacy and security risk they are taking. Unfortunately, past research has shown that neither of these two approaches result in alignment between user preferences, user expectations, and actual privacy settings [Felt et al. 2012].

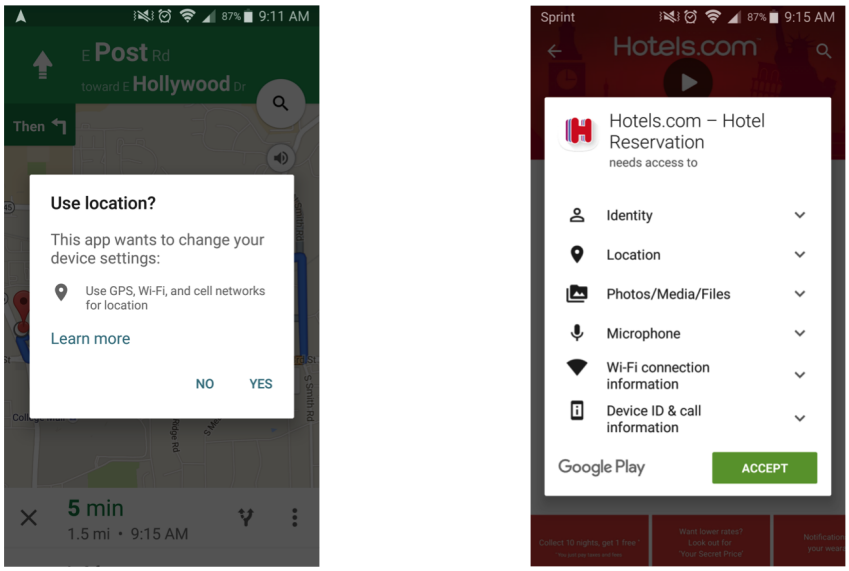

(a) Resource Access Warning (b) Permissions Manifest

Fig. 1: Privacy risk communication mechanisms

In this paper, we use visual indicators along with auditory feedback to communicate the aggregate privacy rating of an application to the end user. Our experiment results show that participants with both visual indicators and auditory feedback made app choices primarily based on the Privacy Rating of the application while the participants in the control group made app choices primarily based on App Rating.

2. BACKGROUND

2.1. Current permissions models

Android and iOS are the two leading operating systems for smartphones and mobile devices. While these two operating systems automatically grant applications permissions to access resources that pose very little risk to users privacy and security, they require users to explicitly grant permission to more sensitive resources. However, the manner in which these permissions requests were presented were previously very different in these operating systems. Android used to present users with install time permissions manifests (shown in Figure 1b). The users had the option to install the application by granting it all the permissions in the manifest or they could deny the permissions and not install the application. This is still the case for devices running Android 5.1 or lower. But for Android 6.0, Google choose to move towards the iOS model.

In the iOS model which now has been adapted by Android versions 6.0 and higher, the users are presented with permissions requests during run time. These permissions requests are presented in the form of warning dialogs and are popularly known as resource access warnings. For example, the first time an application tries to access a users location he/she sees a warning dialog similar to the one in Figure 1a. At this point the user has the option to grant or deny the permission. If the user chooses to grant this permission, then he/she wont be shown the resource access warning the next time the same app tries to access his or her location. Users also have the option to re-

voke these permission by navigating to Privacy Setting in iOS or Application Manager in Android. While iOS's model does give users more freedom with respect to setting custom permissions for each application, it fails to provide users the desired flexibility [Benisch et al. 2011]. Additionally, prior research has also found that the iOS vetting and run-time warnings were less effective than Android's community ratings and permissions manifest mechanism [Han et al. 2014]. Specifically, a comparison of 2600 apps found the iOS versions were consistently over-privileged compared to Android [Han et al. 2013]. Therefore, expecting a change by replacing the Android permissions model with the iOS model to resolve the privacy paradox in the case of smart phones seems unduly optimistic.

2.2. Drawbacks of existing permissions models

Neither of the two permissions models has proven to be successful in providing consumers with actionable information for making informed decisions [Agarwal and Hall 2013]. Therefore, both iOS and Android users are largely unaware of the resources accessed by the apps [Mylonas et al. 2013]. One of the reasons for this is the users' habituation to ignore the current interactions presented in both Android and iOS permissions models. In the case of textual warnings or permissions manifestes used in Android, past research has shown that people usually ignore or pay little attention to them [Felt et al. 2012]. More specifically, a series of online surveys and laboratory studies conducted by Felt et al. found that only 17% of the participants paid attention to permissions during app installation [Felt et al. 2012]. Consumers are also accustomed to ignoring resource access warnings. Warning dialogs are excessively used in todays computers and mobile devices. This overuse of warning dialogs has desensitized people towards them [Anderson et al. 2016; Vance et al. 2017]. Therefore, people view these warning dialogs as interruptions rather than security/privacy alerts and click through them to get on with their current task [Sunshine et al. 2009; Xia and Brustoloni 2005; Egelman et al. 2008; Brustoloni and Villamarín-Salomón 2007].

Users' inability to comprehend the permissions presented to them and their implications is another reason why the current permissions models are unsuccessful. Textual warning in permissions manifests, for example, are commonly requested in English with too much technical jargon which effectively assumes that all smartphone users possess an above average level of basic literacy in addition to computer literacy required to comprehend the permissions information and translate to the risks of agreeing to the requested permissions. However, this is not the case. Not all smartphone users have basic education or computer literacy. As a result, they do not understand the technical jargon used to describe permissions or the implications of providing sensitive permissions to applications [Kelley et al. 2012; Felt et al. 2012].

Therefore, even though people value their online privacy [Nissenbaum 1998], they are unable to make privacy preserving decisions as the current permissions models fails to provide them with actionable risk information. Hence, after considering the above mentioned drawbacks of the existing permissions models, researchers have come up with alternative approaches to provide smartphone users with actionable privacy risk information so that they can make privacy preserving decisions. In this paper, what we mean by "privacy preserving decisions" or "privacy preserving behaviour" is that people/users consider the privacy implications of the permissions requested by an app into consideration when selecting an app.

2.3. Privacy Indicators

As mentioned above, not everyone has the basic education and the computer literacy to understand the information presented in the privacy warning and the risks of giving access to sensitive resources. In such cases, simple privacy indicators that summa-

rize the privacy risks can be beneficial. The use of social cues as privacy indicators was found to be most effective. Eyes were a social cue that were commonly used by researchers to communicate privacy risk. Liccardi et. al. used eyes to communicate sensitivity score (risk score for an app) and highlight risky permissions in Androids permissions manifest [Liccardi et al. 2014]. Schlegel et. al. used eyes appearing on the home screen of a smartphone and growing in size to represent the number of accesses granted to a users location [Schlegel et al. 2011]. Benton et. al and Rajivan et. al. used emoticons in addition to eyes to communicate an aggregate privacy risk to the users. Rajivan et. al. also used padlocks to communicate privacy benefit rather than privacy risk to the end user [Benton et al. 2013; Rajivan and Camp 2016]. So, most of the privacy indicators used by researchers were textual or visual in nature. No one used aural cues as privacy indicators to communicate the risk posed by app.

2.4. Framing of Privacy

Researchers also explored positive and negative framing and how it affected user decisions. Here, positive framing refers to communicating privacy benefit or the privacy offered by an app while negative negative framing refers to communication privacy risk posed by an app. The use of positive framing for improved security and privacy interactions is generally supported by work in the psychology of security, although not consistently applied in the case of apps [Garg and Camp 2013; West 2008; Acquisti et al. 2015b]. One comparison of different icons showed that the positive lock icon was more effective than a negative framing of eyes or sad emoticons [Rajivan and Camp 2016]. In contrast, Cho et. al. found limited efficacy for either, and that there was little significance between positive and negative framing [Choe et al. 2013] . In a follow-up experiment, the authors determined that visual cues could have an effect on participants‘ permissions-based app decisions. That effect was measured by presenting participants with the same app repeatedly, and by asking them to make a comparison between the two scales (negative and positive). Participants were also asked to make choices either based on positive framing or negative framing of risk and found that participants made more risk-averse choices with positive framing in comparison to negative framing [Chen et al. 2015].

2.5. Timing

Timing of permission presentation has also been previously explored. Balebako et. al. investigated the amount of attention users paid to permissions notices when they were presented in the app store, when an app was launched, during app use and after app use. Their results showed that people paid more attention to permissions when they were presented during app use [Balebako et al. 2015]. Their results also showed that users are unlikely to pay attention to permissions shown in the app store. However, in this case, attention was measured by using recall as a proxy and not based on the decision users made or their behaviour. Moreover, it does not matter if people can recall the permissions an app asked for as long as the permissions notices can assist users with their permissions preferences. For example, Kelley et. al. found that when permissions were included in the app description page instead of presenting them after the users chose to install an application users chose applications that had fewer permissions [Kelley et al. 2013]. But, in their study they asked participants to imagine that they were choosing the apps for a friend. From a risk point of view, people are more accurate in their risk estimates when making judgements about acceptability of risk for others. Availability, affect, assimilation and representativeness can all result in different estimates for privacy risk for oneself as compared to a friend [Garg and Camp 2013]. In general, people have been found to be more impartial and risk averse while recommending a risky situation to others [Helfinstein et al. 2015].

2.6. Our Approach

We combined insights from previous work in mobile privacy and in general risk communication to design an alternative mechanism that provided actionable and usable information to people concerned about privacy. In pursuit of this goal, we have augmented the current permissions model with simple icons and sounds. Building on previous work described above, we designed the icons to provide positive beneficial framing for protecting privacy, and aural cues to provide feedback as a form of priming. We then tested four groups: control, icons only, sound only, and icons combined with sound.

Typically, privacy priming is done using a privacy survey, notification, feedback, or other design nudges. Surveys for app installations in the real world are not workable. Nudging and feedback create additional privacy management tasks. Since our design goal was to inform users about an app's privacy without creating additional tasks we did not pursue nudging and feedback. We also did not want to interrupt the app installation process, so we decided against video priming. We rejected haptic communication because vibrations are not clearly good or bad, although it is simple to communicate intensity.

For this experiment we decided to use aural cues. More specifically, we used cheers and jeers. Not only can users easily comprehend the positive nature of joyous cheers and the negative implication of angry jeering without any additional cognitive effort, they are also not interrupted to do additional tasks. Our results, showed that users with both visual and aural cues were more likely to make app choices primarily based on privacy rating.

3. METHOD

3.1. Aural Cues

In this experiment, we explored two different aural cues to prime users for privacy. One of the aural cues was an audio snippet of people cheering. This was played when a participant selected an app with a high privacy rating. We hypothesize that this positive feedback would encourage people to select more apps with a high privacy rating in the future. The other aural cue that was used in this experiment was an audio snippet of people booing. This audio snippet was played when a participant selected an app with a low privacy rating. We hypothesize that this negative feedback would alert people about the privacy rating of the app and prevent them from installing it. Additionally, we hypothesized that this would prevent people from selecting apps with a low privacy rating in the future. The two audio snippets can be found at Cheering, Booing.

Fig. 2: Padlock Icon used to communicate Privacy Rating

Privacy Cues	Group 1: Control	Group 2: Lock	Group 3 : Sound	Group 4: Sound and Lock
Permissions Manifest	yes	yes	yes	yes
Padlock Privacy Rating	no	yes	no	yes
Sound Notifications	no	no	yes	yes

Table I: List of features available in different experimental groups

3.2. Visual Indicator for Privacy Rating

Rajivan et al. studied the effectiveness of of three different visual indicators (frowning face, eye and lock icon) and two different framing mechanisms (positive and negative framing) to communicate privacy risk. The results showed that participants who were presented with positive framing using the padlock made consistently different choices when compared to participants in the control group. Therefore, for our experiment we employed positive framing using the padlock to communicate the aggregate privacy rating.

Since we employed positive framing in our experiment, we communicated privacy benefit rather than privacy risk to the user. For example, an app that posed low risk to users' privacy had a high privacy rating, and vice versa. Another way of looking at this is that more locks imply more privacy. The padlock icon used in our experiment can be found in Figure 2.

3.3. Experimental Groups

To measure the change in behaviour caused by the privacy cues, we divided our participants into four groups: Control Group, Lock Group, Sound Group and Lock and Sound Group. The participants of the Control Group were provided with a version of the interactive simulator that was an exact simulation of Googles Play Store. In other words, the participants in the Control Group were not provided with the aggregate privacy ratings, but they had access to the permissions manifest. The participants in Lock Group were provided with the aggregate privacy rating for each and every app using the visual indicator proposed in section 3.2. The participants in the Sound Group heard cheers/jeers based on the app's privacy rating but were not provided with a visual representation of the aggregate privacy rating. Finally, the participants in the Lock and Sound groups were provided with both visual and aural cues. Table I shows privacy cues that were available to the participants in each these groups.

3.4. Interactive Pay Store Simulator

As the primary purpose of our experiment was to investigate if the proposed visual and aural cues are effective at communicating the privacy rating to the end user, it was important for us to simulate a realistic app installation environment which would trigger some of the same cognitive processes involved in real world app installations. In order to do so, we built an interactive play store simulator which simulated Googles Play Store. The simulator which ran on a web browser simulated two primary aspects of the Play Store:

(a) The *list of apps page*: Android allows its users to browse applications by category. So when a user selects a category of interest they can see a list of apps in that category. As you can see in Figure 3a, our simulator simulates the window that displays a list of apps in a particular category.

(b) The *app description page*: when a user selects an app in the Play Store by clicking on it, he/she is redirected to the app description page. The app description page on the

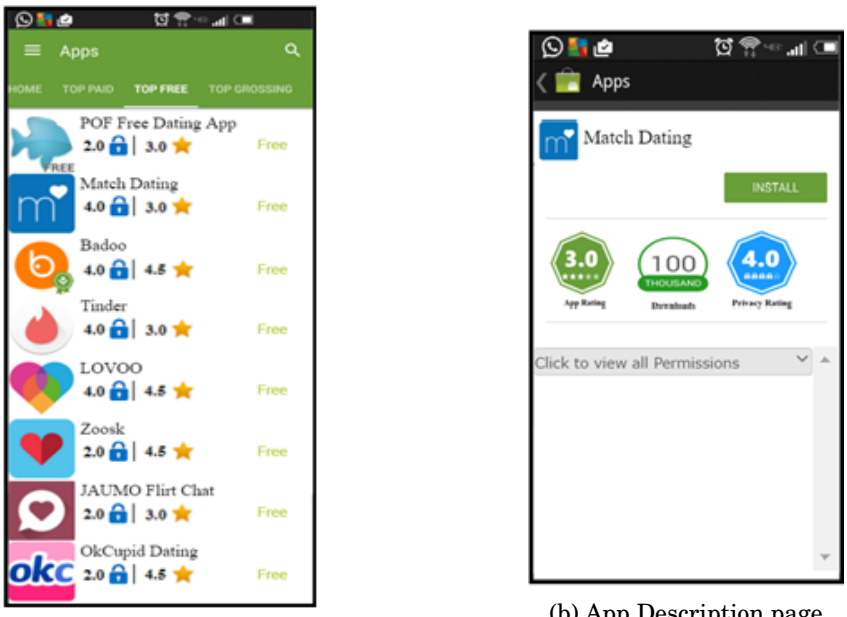

(a) List of Apps page

(b) App Description page

Fig. 3: Screenshots of the simulated Android app store interfaces.

Play Store provides users with app rating, download count and a permissions manifest. A simulation of this page can in found in Figure 3b.

In addition to the information provided by the Play Store on the *list of apps* and the *app description* pages, we provide participants in all the experimental groups with privacy rating. As described in section 3.3 , the format in which Privacy Rating is provided differs from group to group. In the Lock Group, Privacy Rating is provided using the visual cues. In the Sound Group, privacy rating is provided using aural cues. In the Lock and Sound group, privacy rating is provided using both visual and aural cues. As you can see in Figure 3, the visual cues for privacy rating were provided alongside the app rating. The aural cues on the other hand were played when a participant selected an app in the *list of apps* page. An illustration of this can be found in Figure 4

Analogous to Androids Play Store, participants were able to move back and forth between the above mentioned pages using the back arrow, install and uninstall applications and view the permissions manifest by clicking on the *click to view all permissions* drop down.

3.5. Apps

The goal of this experiment was not only to test the efficacy of cues, but also to further investigate the applicability of the results in a practical interaction design. Therefore, we decided to use popular apps in our application. This decision to use popular apps was motivated by research indicating the importance of popularity and downloads in decision-making beyond the general importance of familiarity and the perceptions of predictability in decision-making [**?**]. A series of surveys, interviews and focus groups illustrated that popularity indicates acceptability of privacy policies, with use by others being an implicit, environmental cue [Morton 2014]. Extensive investigation of risk perception in mobile and wearable devices found that familiar technologies were per-

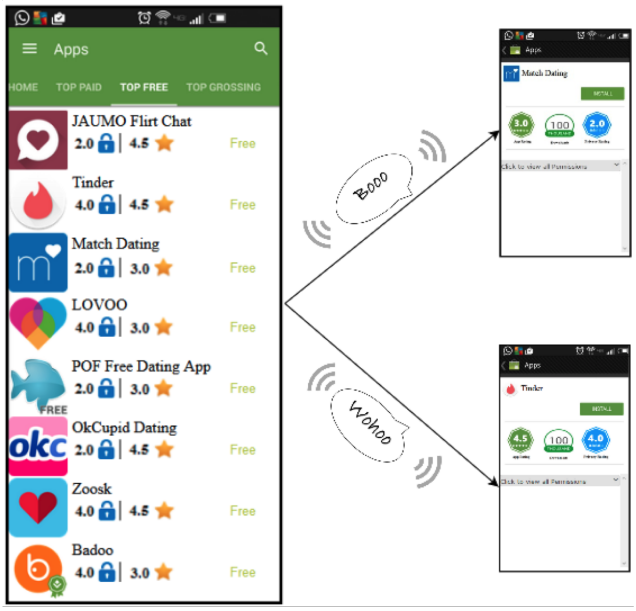

Fig. 4: Sounds associated with apps. Apps with a privacy rating of four or five result in cheers and those with a privacy rating of one or two result in jeers.

ceived as less risky [Lee et al. 2015]. More popular apps are by definition more likely to be familiar. Download counts and popularity are both indicators of reputation. Multi-disciplinary investigations of trust online have shown reputation to be a critical factor in decision-making [Anthony et al. 2010]. An investigation of the analyses of trust in different domains resulted in identification of common factors. Familiarity and reputation are consistently factors in trust decisions in a wide range on online environments [Costante et al. 2015].

3.6. Experimental Variables

Prior research has shown that permissions manifests are ineffective at communicating privacy risk [Benton et al. 2013]. So people rely on other attributes, like Download Count and App Rating of mobile applications to make their app choices [Benton et al. 2013; Rajivan and Camp 2016]. Therefore, we expected control group participants to install applications with high app rating and download count. On the other hand, since participants in the remaining experimental groups were presented with privacy rating in a more comprehensible fashion, we expected them to install applications with a higher privacy rating.

In order to determine the validity of our above mentioned hypothesis, we recorded three dependent variables for each and every app installed by the participants in all four experimental groups. These dependent variables are: App Rating, Privacy Rating and Download Count. App Rating is the combined user rating of an app and Download Count is number of times the app was downloaded from the Play store. While Android currently uses permissions manifests to communicate privacy risk to the user at install time, we use the Privacy Rating metric to communicate risk to the end user. It is conveyed to the end user through visual cues or aural cues or both visual and aural cues (depending on the experiment group 3.3). Ideally Privacy Rating would be com-

puted based on the information accessed and leaked by the application. Here we used a Wizard of Oz system and Privacy Ratings were assigned by experimenters.

3.7. Experimental Procedure

The participants for this study were recruited from Amazons Mechanical Turk (MTurk). The experiment was posted as an HIT (Human Intelligent Task) on MTurk. The first page of the HIT displayed information about the experiment. MTurk users could review this information and decide on whether to participate in the experiment or not. Upon agreeing to participate in the in the study, all the participants were provided with a simple set of instructions on how to use the interactive Play Store simulator. After reading the instructions, the participants were allowed to move on to the simulated environment and start making application choices. They were presented with 2 sets of application categories with 8 applications in each category. The order of categories and the order of applications under each category was randomized for all participants.

Participants were asked to make at least 4 application choices in the order of their preference (1st choice being the most preferred and the 4th choice being the least preferred) for each category. Once the participants made all the necessary application choices, they were presented with a set of questionnaires to understand their app installation behavior and capture their computer literacy and demographics.

4. RESULTS

Before we proceed to analysis we summarize our results. The study has 4 experimental conditions. 80 participants were recruited for each experimental condition. In total, we enrolled 320 participants for our study.

4.1. Data Collected

Apart from the app choices and the responses to the questionnaires, we also collected several implicit data measures from the experiment which include permissions viewed, amount of time spent on choosing apps in each category and the total time the participants took to complete the experiment.

4.2. Exclusion Criteria

As mentioned earlier, we recruited 80 participants for each group and a total of 320 participants took part in the experiment. Out of the 320 participants, 17 participants were disqualified for providing contradicting answers to questions in the questionnaires. For example, the question “Do you review/read the permissions presented to you before you install an application from Android Play Store?” was asked twice and if a participant gave two different answers then he or she was disqualified. We also excluded all the results from the participants who took less than 3 minutes to complete the Study. After applying the above mentioned exclusion criteria, we ended up with a total of 235 participants. These exclusion criteria was used to identify participants who did not put any cognitive effort towards making app choices.

4.3. Demographics

All the participants for this study were recruited from within the United States. This was done by applying the location qualification in MTurk. Out of the 235 participants, 60.85% were male and 39.15% were female. The average age of the participants was 31 years.

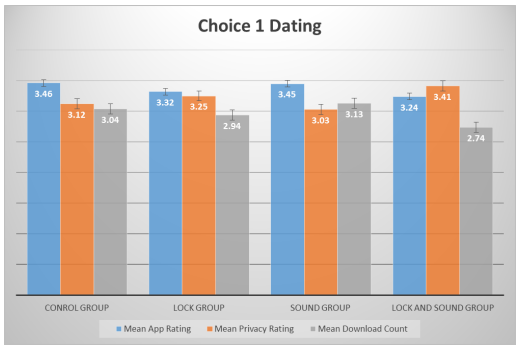

(a) Means for App Choice 1

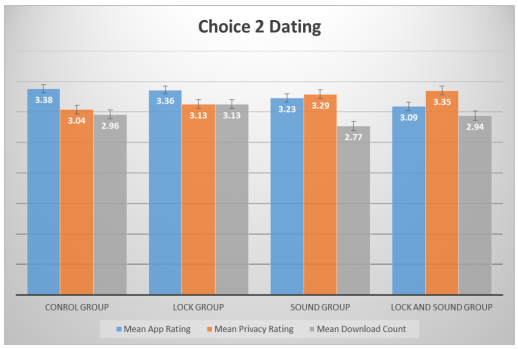

(b) Means for App Choice 2

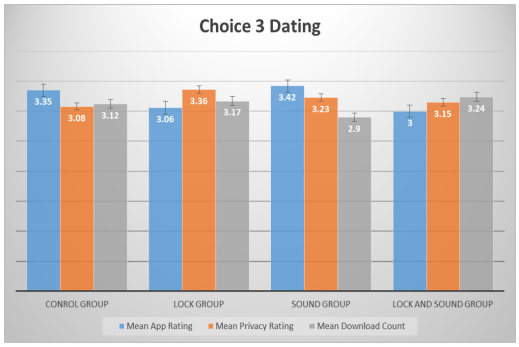

(c) Means for App Choice 3

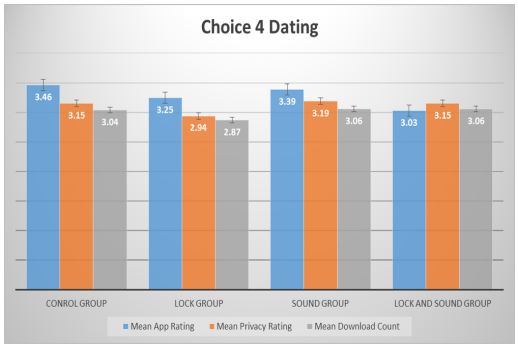

(d) Means for App Choice 4

Fig. 5: Mean Values for Dating Apps

4.4. Data Description

4.4.1. Basic Means Comparison. Figure 5 has a collection of four histograms with mean App Rating, Privacy Rating and Download Count for all four app choices in the dating category. As you can see in Figure 5, the mean app rating for all four choices in the control group is higher than the mean Privacy Rating and the mean Download Count. This indicates that app rating had a higher influence on the participants app choices in the Control Group when compared to privacy rating and download count. Also, the mean Privacy Rating is always higher than the mean App Rating and Download Count in the Lock and Sound Group with choice 3 being the only exception (Mean Download Count (3.24) is greater than Mean Privacy Rating (3.15)). This implies that Privacy Rating had the most influence on the participants app choices in the Lock and Sound Group. The mean Privacy Score of the Lock and Sound Group is higher than the mean Privacy Rating of the Control Group for the first three app choices. The mean Privacy Ratings for the 4^{th} app choice are the same for both groups. The Lock Group and the Sound Group also consistently had a higher mean Privacy Rating when compared to the Control Group with choice 1 being an exception for the Sound Group (Control Group (3.12) $>$ Sound Group (3.03)) and choice 4 being an exception for the Lock Group (Control Group (3.15) $>$ Sound Group (2.94)). This implies that the influence of Privacy Rating was higher when the participants were provided with the privacy cues. This behavior is very prominent in the Lock and Sound Group.

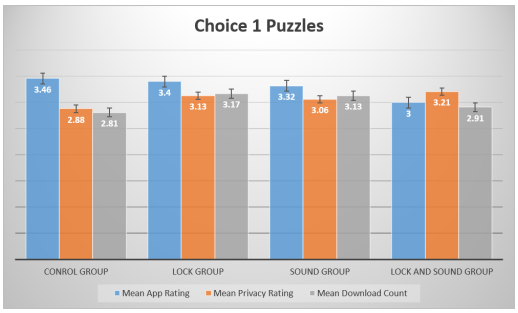

(a) Means for App Choice 1

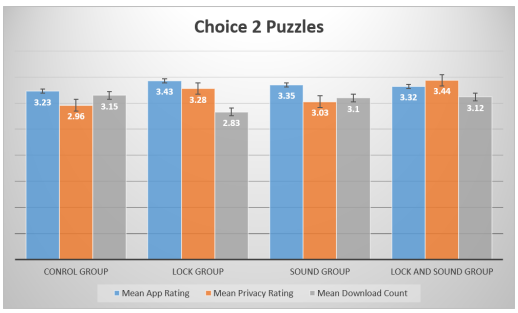

(b) Means for App Choice 2

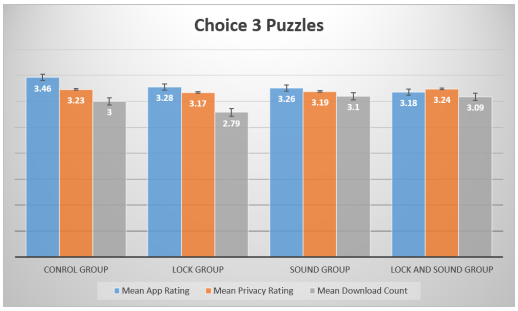

(c) Means for App Choice 3

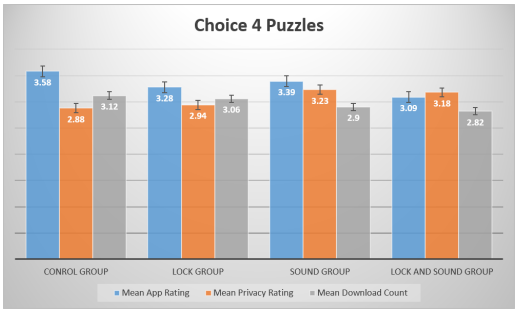

(d) Means for App Choice 4

Fig. 6: Mean Values for Puzzle Apps

Figure 6 has a collection of four histograms with mean App Rating, Privacy Rating and Download Count for all four app choices in the Puzzles category. Similar to what we saw for the dating apps, the mean app rating for all four app choices in the control group is higher than the mean privacy rating and download count indicating that participants in the control group made their app choices primarily based on app rating. One other similarity that we observed was that the mean Privacy rating for all four choices in the Lock and Sound Group higher than the mean App Rating and Download Count. This indicates that Privacy Rating has a higher influence on the app choices made by the participants in the Lock and Sound group when compared to the Control Group. The prior indication is further strengthened by the fact that the mean Privacy Rating for the Lock and Sound group is higher than that of the Control Group for all four app choices. Also similar to the Dating Apps, the mean Privacy Rating for the Lock Group and the Sound Group is higher than that of the Control Group for 3 out of 4 app choices (mean Privacy Rating Control group $(3.23) >$ mean Privacy Rating Sound Group $(3.19) >$ mean Privacy Rating Lock Group (3.17) for Choice 3). This reaffirms our belief that Privacy Rating had a higher influence on app choices made by participants in groups with privacy cues when compared to the Control Group. Once again, this trend is more prominent in the Lock and Sound Group.

4.4.2. App Installation Frequency. Please recall that at the end of the survey we asked participants we asked participants how often they installed applications from Google's Play Store. As you can see in Figure 7, the response to the question is as follows:

— 0% of the participants indicated that they *never* installed apps from the app store.

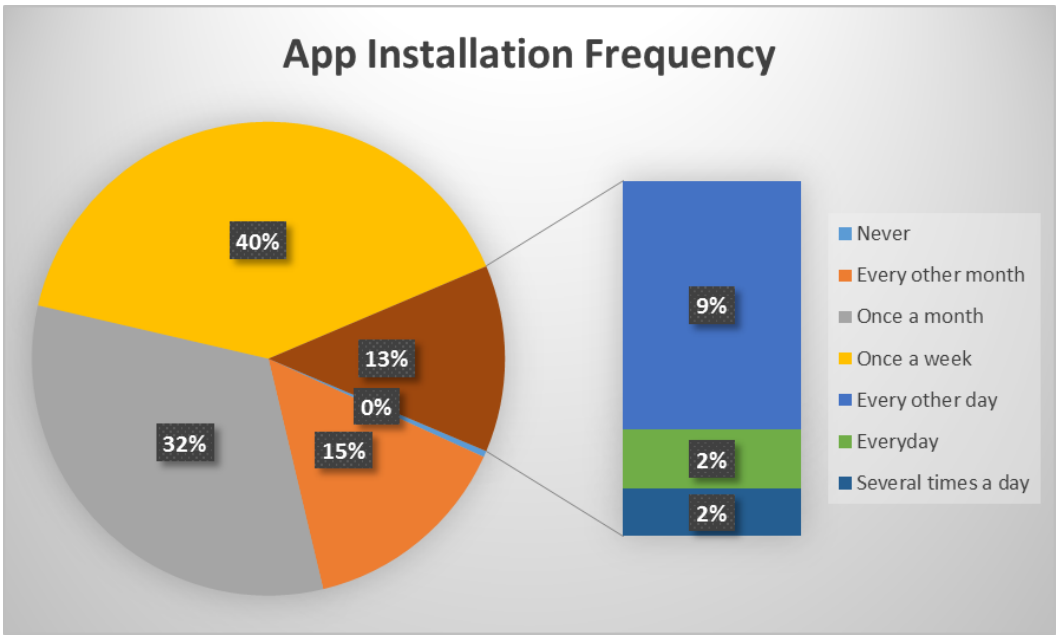

Fig. 7: Self Reported App Installation Frequncies of Participants in The Study

— 15% of the participants reported that they installed apps *every other month*.
— 32% of the participants reported that they installed apps *once a month*.
— 40% of the participants reported that they installed apps *once a week*.
— 9% of the participants reported that they installed apps *every other day*.
— 2% of the participants reported that they installed apps *everyday*.
— 2% of the participants reported that they installed apps *several times in a day*.

In summary, since 87% of the participants reported that they installed apps from the Play Store *once a week* or *less than once a week*, this indicates that app installation is an activity that does not occur very often.

4.5. Analysis

Please recall that our experiment has one Control Group and three experimental groups. The participants in the Control Group interacted with a play store simulator that was an exact replication of Googles Play Store while the participants in the experimental groups received additional cues related to the apps privacy. The three experimental groups are:

— Lock Group Privacy Rating of the App was communicated using visual cues.
— Sound Group Privacy Rating of the App was communicated using aural cues.
— Lock and Sound Group Privacy rating of the App was communicated using both aural and visual cues.

Participants in all four groups were asked to install four out of eight apps for each category (Dating and Puzzles). Participants were also asked to install these apps in the order of their preference (the first app installation should be the most preferred and the last app installation being the least preferred). For each app installed by the participants, we recorded three dependent variables: app rating, privacy rating and download count. These three dependent variables were normalized to be in the same range

						95% Confidence Interval For Exp(B)	
		p-values	B	Exp(B)	1/Exp(B)	Lower	Higher
Lock and Sound Group	Dating Apps	p <0.001	-0.642	0.526	1.9	0.392	0.707
	Puzzle Apps	p <0.001	-0.731	0.481	2.07	0.353	0.659
Lock Group	Dating Apps	0.048	-0.326	0.722	1.38	0.522	0.997
	Puzzle Apps	0.052	-0.311	0.733	1.36	0.536	1.003
Sound Group	Dating Apps	0.107	-0.247	0.781	1.28	0.579	1.055
	Puzzle Apps	0.032	-0.300	0.741	1.35	0.563	0.975

Table II: GEE results for PrivacyOverAppRating

for ease of analysis. Additionally, we computed two more dependent variables: PrivacyOverAppRating and PrivacyOverDownloadCount. PrivacyOverAppRating measures if the Privacy Rating for an app is higher (1) or lower (-1) or the same (0) as the App Rating. Similarly, PrivacyOverDownloadCount measures if the Privacy Rating for an app is higher (1) or lower (-1) or the same (0) as the Download Count.

We hypothesized that people with privacy cues are more likely to make app choices with a higher privacy rating when compared to people with no privacy cues. So we use Generalized Estimation Equations to test our hypothesis.

4.6. Generalized Estimation Equations

Generalize Estimating Equations (GEE) are an extension of Generalized Linear Models and are commonly used to analyze correlated data that arises from repeated measurements. In our case, the repeated measurements stem from each participant making four app installations in each category. The specific goal of changing the interaction is to create systematic change in the selection of apps in order to reduce the overall risk of information exhilaration. A GEE analysis can evaluate the aggregate decisions to see if users in different groups behaved differently as a whole. GEE does not restrict the dependent variables to be continuous or have a normal distribution. GEE aligns with our experimental goals and the resulting data.

4.6.1. Privacy vs Other Dependent Variables. We were interested in measuring the effect of simple privacy cues on app installation choices. More specifically, we wanted to understand if participants with privacy cues installed apps with a higher Privacy Rating. Hence, we computed the analysis on the dependent variables PrivacyOverAppRating and PrivacyOverDownloadCount, as these variables record comparisons of privacy rating against App Rating and Download Count respectively. To be more descriptive, the PrivacyOverAppRating tells us if Privacy Rating for an installed app is higher than, equal to or lower than its App Rating. Similarly, PrivacyOverDownloadCount tells us if Privacy Rating for an installed app is higher than, equal to or lower than its Download count. The results from the data analysis for PrivacyOverAppRating and PrivacyOverDownloadCount can be found in Tables II and III respectively.

As you can see in Table II, PrivacyOverAppRating is statistically different between the Control Group and the Lock and Sound Group for both Dating and Puzzle apps. The coefficient estimate B and the odds ratio Exp(B) can also be found in the Table II. It is usually the latter that are more informative. In this case, the odds of the Control Group are 0.526 times that of the Lock and Sound Group to have a higher value in the response variable(PrivacyOverAppRating) for dating apps. Since the odds are less than one we compute 1/odds (1.9) to make the interpretation more clear and understandable. Therefore, the odds of the Lock and Sound Group having a higher value is 1.9 times that of the Control Group for dating apps. Similarly, the odds of the

						95% Confidence Interval For Exp(B)	
		p-values	B	Exp(B)	1/Exp(B)	Lower	Higher
Lock and Sound Group	Dating Apps	0.033	-0.285	0.752	1.33	0.579	0.977
	Puzzle Apps	0.001	-0.442	0.643	1.55	0.497	0.833
Lock Group	Dating Apps	0.406	-0.110	0.896	1.11	0.692	1.160
	Puzzle Apps	0.031	-0.273	0.761	1.31	0.594	0.975
Sound Group	Dating Apps	0.064	-0.238	0.788	1.26	0.613	1.014
	Puzzle Apps	0.196	-0.152	0.859	1.16	0.683	1.081

Table III: GEE results for PrivacyOverDownloadCount

						95% Confidence Interval For Exp(B)	
		p-values	B	Exp(B)	1/Exp(B)	Lower	Higher
Lock and Sound Group	Dating Apps	0.028	0.352	-0.704	1.420	0.514	0.963
	Puzzle Apps	$p < 0.001$	0.566	-0.568	1.760	0.418	0.771
Lock Group	Dating Apps	0.307	0.152	-0.859	1.164	0.641	1.150
	Puzzle Apps	0.047	0.288	-0.750	1.334	0.565	0.996
Sound Group	Dating Apps	0.139	0.183	-0.832	1.201	0.653	1.061
	Puzzle Apps	0.036	0.280	-0.756	1.323	0.582	0.981

Table IV: GEE results for Privacy Rating

of the Lock and Sound Group having a higher value is 2.07 times that of the Control Group for puzzle apps.

In Table III, you can see that PrivacyOverDownloadCount is also statistically significant between the Control Group and the Lock and Sound Group for both dating and puzzle apps. The odds of the Lock and Sound Group having a higher value for PrivacyOverDownloadCount is 1.33 (1/Exp(B)) times that of the Control Group for Dating Apps. The odds of the Lock and Sound Group having a higher value for PrivacyOverDownloadCount is 1.55 (1/Exp(B)) times that of the Control Group for Puzzle Apps.

Therefore, the results show that the Lock and Sound Group is more likely to have a higher value for both PrivacyOverAppRating and PrivacyOverDownloadCount. Please recall that a higher value for PrivacyOverAppRating or PrivacyOverDownloadCount implies that the Privacy Rating is higher than App Rating or Download Count respectively. Hence, it can be inferred that the participants with both visual and aural cues are more likely to make their app installation choices primarily based on Privacy Rating.

The results were not as significant for the Lock Group and the Sound Group. For the Lock Group, the PrivacyOverAppRating was found to be statistically different only for Dating apps and PrivacyOverDownloadCount was found to be statistically different only for Puzzle Apps. For the Sound Group, PrivacyOverAppRating was found to be statistically different only for Puzzle apps and PrivacyOverDownloadCount was not statistically different for either Puzzle or Dating Apps. The p-values along with the odds ratios can be found in Table II and III.

4.6.2. Privacy Rating, App Rating and Download Count. We also performed GEE analysis on the data to see if Privacy Rating, App Rating and Download Count were statistically significant between the Control Group and the Experimental Groups. Our results show that Privacy Rating and App Rating are statistically significant between the Control Group and the Lock and Sound Group for both dating and puzzle apps. For the Lock Group and the Sound Group, Privacy rating is statistically different only for the puzzle

					95% Confidence Interval For Exp(B)	
		p-values	B	Exp(B)	Lower	Higher
Lock and Sound Group	Dating Apps	p <0.001	0.705	2.025	1.441	2.845
	Puzzle Apps	0.001	0.631	1.88	1.303	2.712
Lock Group	Dating Apps	0.066	0.381	1.464	0.975	2.197
	Puzzle Apps	0.348	0.198	1.219	0.806	1.844
Sound Group	Dating Apps	0.620	0.101	1.106	0.742	1.648
	Puzzle Apps	0.195	0.242	1.247	0.883	1.838

Table V: GEE results for App Rating

apps and App Rating is not statistically different from the Control Group for both puzzle and dating apps. Finally, Download Count is not statistically significant between the Control Group and the Experimental Groups for both dating and puzzle apps. The GEE analysis results for Privacy Rating and Download Count can be found in Table IV and Table V.

Upon further examination, we found that:

— The odds of participants in the Lock and Sound Group having a higher value for Privacy Rating is 1.42 (1/Exp(B)) times that of the Control Group for Dating Apps. Similarly, for puzzle apps, the odds of participants in the Lock and Sound Group having a higher value for Privacy Rating is 1.76 (1/Exp(B)) times that of the Control Group.
— The odds of participants in the Lock Group or the Sound Group having a higher value for Privacy Rating is 1.3 times (1/Exp(B)) times that of the Control Group for puzzle apps.
— The odds of participants in the Control Group having a higher value for App Rating is 1.44 (Exp(B)) times that of the Lock and Sound Group for dating apps. Similarly, for the puzzle apps, the odd are 1.88 times that of the Lock and Sound Group.

Therefore, a one on one comparison between the Control Group and the Lock and Sound Group shows that participants in the Control Group made app choices primarily based on App Rating while participants in the Lock and Sound Group made app choices primarily based on Privacy Rating.

4.7. A Note on Analysis

A standard approach to analysis of this type of data is to implement a two sided t-test for each pair of selections and then determine if these are significant. Reviewers may expect a standard p-value to determine if the results are significant. Typically, to determine if the difference between Groups mentioned above had any statistical significance researcher would perform One Way ANOVA and post-hoc tukey (pairwise comparison) tests as they are commonly used to determine statistical differences between groups. Therefore, we include this section to note both these results, and to explain why these are not the results we are looking for. Essentially, a two-sided t-test comparing the selections of each group was similar to the results reported above with sound and lock being consistently significant, and sound or lock showing significance in to choices. However, in order for One Way ANOVA or post-hoc tukey tests to generate accurate results the study data must meet certain assumptions. These assumptions are as follows:

— The dependent variable must be measured on a continuous scale.
— The independent variable (in our case Groups) should have two are more categories.
— The observations must be independent (i.e. there should be no relationship between observations in each Group or between Groups).
— The data should not have significant number of outliers.
— The dependent variable must be normally distributed.

Our study violates two out of the five assumptions. Our dependent variables are not recorded on a continuous scale. App Rating, Privacy Rating and Download Count are nominal (take values 2 or 4). PrivacyOverAppRating and PivacyOverDownloadCount are multinomial (take values -1 or 0 or 1). Since our dependent variables are not recorded on a continuous scale they are also not normally distributed. In fact, Shapiro-Wilks normality test reveals that none of the dependent variables are normally distributed ($p < 0.001$).Thus, we do not include these tests of significance. Therefore, in the quest to more accurately determine the statistical differences between the Control Group and the Experimental Groups we used Generalized Estimation Equations.

5. DISCUSSION

As mentioned in Section 4.5, the results clearly show that people with both visual indicators and aural cues are more likely to select apps with a higher Privacy Rating. In fact, people with both privacy cues were more likely to choose apps that whose Privacy Rating is higher than both App Rating and Download Count. This indicates that when people are presented with both privacy cues they primarily make their app choices based on Privacy Rating.

For participants who were presented with just visual indicators or just aural cues, Privacy Rating was not statistically different from that of the Control Group for dating apps. But the results do show that participants with just one of the two privacy indicators/cues were more likely to make app choices with a higher Privacy Rating for puzzle apps.

One reason for this disparity between the app choices for dating and puzzle apps could be that the participants were more familiar with popular dating apps when compared to popular puzzle apps. Since familiar technologies are perceived to be less risky [Lee et al. 2015], it could be possible that participants considered familiar dating apps to be less risky even though the Privacy Rating suggested otherwise. Another reason for this could be that participants were more willing to share sensitive information with dating apps when compared to puzzle apps. Dating apps require people share a lot of sensitive resources like their location, camera etc. But people can easily justify why a dating app would require access to these sensitive resources. For example, it is easy to comprehend that a dating app requires access to users' location to find people around them. But the same cannot be said about puzzle apps or games apps in general. For example, in a study conducted by Shklovski et. al. participants felt deceived and expressed severe concerns when they learnt about data collected by the fruit ninja app [Shklovski et al. 2014]. Whatever the case may be, it is clear that people with just one privacy cue do not always take privacy of an app into consideration.

Even if Privacy Rating is provide alongside App Rating people may not always take it into consideration while selecting an app to install (Similar to what we saw for dating apps). So without an audio feedback to alert the users about the low privacy rating users are unlikely to pay attention to it. On the other hand, if people are provide with audio feedback but no visual indicator for privacy rating then they may not be able to understand what the feedback indicates. this lack of comprehension could lead to people selecting apps with a low Privacy Rating.

Finally, regardless of cues, download count information was not significant in the app decision making process. Part of the reason could be that the download count values used for the experiment (50,000 versus 100,000 downloads) were not sufficiently different to have an influence on app choices. The other reason could be that findings which indicate that download count dominates decision processes may have been observing a hidden variable (for example, order of presentation).

Excessive use of warning dialogs has desensitize people towards them [Anderson et al. 2016; Vance et al. 2017]. As audio feedback proposed in this paper is also a form of warning there is a possibility that its excessive use over time might cause people to ignore them. But the results from our survey indicates that app installation is an activity that does not occur very often. As 87% of our participants reported that they installed apps *once a week* or *less than once a week*. Also, unlike warning dialogs our audio feedback alert is unique and is not used by other computing devices. Hence, we believe that the possibility of users getting desensitized to our audio feedback is very low. Nevertheless, a longitudinal study needs to be done to better understand the effects of our visual indicators and aural cues on users app selection behavior.

6. CONCLUSION

Our experiment tested the efficacy of the padlocks in the presence of a realistic distribution of apps and a feasible form of priming. While we considered other options like haptic interactions and additional visual framing for priming users for privacy we decided to to with audio feedback. This is because haptic interactions are not clearly good or bad and additional visual framing could be confounding.

The results from our experiment showed that when participants were presented with both visual indicators (using Padlock icon) and aural cues they made risk based app choices i.e. individuals choose apps with a higher Privacy Rating over apps with a higher App Rating. This was significant change in behaviour when compared to the Control Group where participants made app decisions primarily based on App Rating. Hence, the inclusion of immediate ratings and multimedia priming for privacy offers promise for supporting more informed decision-making in online app stores.

Future work includes longitudinal investigations to determine if these effects are a result of familiarity or overall decision support. That is, would an augmented Play Store increase overall privacy and empower privacy-aware decision-making or would people simply become acclimated to choosing low privacy apps [Anderson et al. 2013]. That is, did these function as warnings to which users would become habituated or were they able to provide decision support which would remain valuable?

ACKNOWLEDGMENTS

REFERENCES

Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015a. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.

Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015b. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.

Yuvraj Agarwal and Malcolm Hall. 2013. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 97–110.

Bonnie Brinton Anderson, Jeffrey L. Jenkins, Anthony Vance, C. Brock Kirwan, and David Eargle. 2016. Your memory is working against you: How eye tracking and memory explain habituation to security warnings. *Decision Support Systems* 92 (2016), 3 – 13. DOI:http://dx.doi.org/https://doi.org/10.1016/j.dss.2016.09.010 A Comprehensive Perspective on Information Systems Security - Technical Advances and Behavioral Issues.

B. B. Anderson, B. Kirwan, journal Jenkins, D. Eargle, S. Howard, and A. Vance. 2013. How Polymorphic Warnings Reduce Habituation in the Brain: Insights from an fMRI Study. In *Proceedings of CHI 2015*. ACM. http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6529497

Denise Anthony, James A Kitts, Christopher Masone, and Sean W Smith. 2010. Internet exchange and forms of trust. *Trust and Technology in a Ubiquitous Modern Environment: Theoretical and Methodological Perspectives. Hershey, PA: IGI Global* (2010), 257–269.

Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick McDaniel. 2014. Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. In *ACM SIGPLAN Notices*, Vol. 49. ACM, 259–269.

Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. 2015. The Impact of Timing on the Salience of Smartphone App Privacy Notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM '15)*. ACM, New York, NY, USA, 63–74. DOI:http://dx.doi.org/10.1145/2808117.2808119

Pawan Bali. 2017. Fake BHIM apps confuse people. (2017).

Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. 2011. Capturing Location-privacy Preferences: Quantifying Accuracy and User-burden Tradeoffs. *Personal Ubiquitous Comput.* 15, 7 (Oct. 2011), 679–694. DOI:http://dx.doi.org/10.1007/s00779-010-0346-0

Kevin Benton, L Jean Camp, and Vaibhav Garg. 2013. Studying the effectiveness of android application permissions requests. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2013 IEEE International Conference on*. IEEE, 291–296.

Alastair R Beresford, Andrew Rice, Nicholas Skehin, and Ripduman Sohan. 2011. Mockdroid: trading privacy for application functionality on smartphones. In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*. ACM, 49–54.

José Carlos Brustoloni and Ricardo Villamarín-Salomón. 2007. Improving Security Decisions with Polymorphic and Audited Dialogs. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS '07)*. ACM, New York, NY, USA, 76–85. DOI:http://dx.doi.org/10.1145/1280680.1280691

Jing Chen, Christopher S Gates, Ninghui Li, and Robert W Proctor. 2015. Influence of risk/safety information framing on android app-installation decisions. *Journal of Cognitive Engineering and Decision Making* 9, 2 (2015), 149–168.

Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. 2013. Nudging people away from privacy-invasive mobile apps through visual framing. In *Human-Computer Interaction–INTERACT 2013*. Springer, 74–91.

Elisa Costante, Jerry Hartog, and Milan Petković. 2015. Understanding perceived trust to reduce regret. *Computational Intelligence* 31, 2 (2015), 327–347.

Manuel Egele, Christopher Kruegel, Engin Kirda, and Giovanni Vigna. 2011. PiOS: Detecting Privacy Leaks in iOS Applications.. In *NDSS*.

Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You'Ve Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1065–1074. DOI:http://dx.doi.org/10.1145/1357054.1357219

William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. 2014. TaintDroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)* 32, 2 (2014), 5.

William Enck, Damien Octeau, Patrick McDaniel, and Swarat Chaudhuri. 2011. A Study of Android Application Security.. In *USENIX security symposium*, Vol. 2. 2.

Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android Permissions: User Attention, Comprehension, and Behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12)*. ACM, New York, NY, USA, Article 3, 14 pages. DOI:http://dx.doi.org/10.1145/2335356.2335360

Vaibhav Garg and Jean Camp. 2013. Heuristics and biases: implications for security design. *IEEE Technology and Society Magazine* 32, 1 (2013), 73–79.

Vindu Goel. 2016. Beware, iPhone Users: Fake Retail Apps are Surging Before Holidays. (2016).

Jin Han, Qiang Yan, Debin Gao, Jianying Zhou, and Huijie Robert DENG. 2014. Android or iOS for better privacy protection? (2014).

Jin Han, Qiang Yan, Debin Gao, Jianying Zhou, and Robert H Deng. 2013. Comparing mobile privacy protection through cross-platform applications. (2013).

Sarah M Helfinstein, Jeanette A Mumford, and Russell A Poldrack. 2015. If all your friends jumped off a bridge: The effect of others actions on engagement in and recommendation of risky behaviors. *Journal of experimental psychology: general* 144, 1 (2015), 12.

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

Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3393–3402.

Linda Lee, Serge Egelman, Joong Hwa Lee, and David Wagner. 2015. Risk Perceptions for Wearable Devices. *arXiv preprint arXiv:1504.05694* (2015).

Ilaria Liccardi, Joseph Pato, Daniel J. Weitzner, Hal Abelson, and David De Roure. 2014. No Technical Understanding Required: Helping Users Make Informed Choices About Access to Their Personal Data. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS '14)*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 140–150. DOI:http://dx.doi.org/10.4108/icst.mobiquitous.2014.258066

Philip Marquardt, Arunabh Verma, Henry Carter, and Patrick Traynor. 2011. (Sp)iPhone: Decoding Vibrations from Nearby Keyboards Using Mobile Phone Accelerometers. In *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11)*. ACM, New York, NY, USA, 551–562. DOI:http://dx.doi.org/10.1145/2046707.2046771

Anthony Morton. 2014. All my mates have got it, so it must be okay: Constructing a Richer Understanding of Privacy ConcernsAn Exploratory Focus Group Study. In *Reloading Data Protection*. Springer, 259–298.

Alexios Mylonas, Anastasia Kastania, and Dimitris Gritzalis. 2013. Delegate the smartphone user? Security awareness in smartphone platforms. *Computers Security* 34 (2013), 47 – 66. DOI:http://dx.doi.org/https://doi.org/10.1016/j.cose.2012.11.004

Helen Nissenbaum. 1998. Protecting privacy in an information age: The problem of privacy in public. *Law and philosophy* 17, 5 (1998), 559–596.

Emmanuel Owusu, Jun Han, Sauvik Das, Adrian Perrig, and Joy Zhang. 2012. ACCessory: Password Inference Using Accelerometers on Smartphones. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications (HotMobile '12)*. ACM, New York, NY, USA, Article 9, 6 pages. DOI:http://dx.doi.org/10.1145/2162081.2162095

Prashanth Rajivan and Jean Camp. 2016. Influence of Privacy Attitude and Privacy Cue Framing on Android App Choices. In *Authentication Workshop of the 12th Symposium on Usable Privacy and Security*. USENIX Association.

Roman Schlegel, Apu Kapadia, and Adam J Lee. 2011. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 14.

Irina Shklovski, Scott D Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2347–2356.

Joshua Sunshine, Serge Egelman, Hazim Almuhimedi, Neha Atri, and Lorrie Faith Cranor. 2009. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *Proceedings of the 18th Conference on USENIX Security Symposium (SSYM'09)*. USENIX Association, Berkeley, CA, USA, 399–416. http://dl.acm.org/citation.cfm?id=1855768.1855793

Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.

Anthony Vance, Brock Kirwan, Daniel Bjornn, Jeffrey Jenkins, and Bonnie Brinton Anderson. 2017. What Do We Really Know About How Habituation to Warnings Occurs Over Time?: A Longitudinal fMRI Study of Habituation and Polymorphic Warnings. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 2215–2227. DOI:http://dx.doi.org/10.1145/3025453.3025896

Ryan West. 2008. The psychology of security. *Commun. ACM* 51, 4 (2008), 34–40.

Haidong Xia and José Carlos Brustoloni. 2005. Hardening Web Browsers Against Man-in-the-middle and Eavesdropping Attacks. In *Proceedings of the 14th International Conference on World Wide Web (WWW '05)*. ACM, New York, NY, USA, 489–498. DOI:http://dx.doi.org/10.1145/1060745.1060817

Yajin Zhou, Xinwen Zhang, Xuxian Jiang, and Vincent W Freeh. 2011. Taming information-stealing smartphone applications (on android). In *Trust and Trustworthy Computing*. Springer, 93–107.

Received February 2007; revised March 2009; accepted June 2009