

A Multilateral Privacy Impact Analysis Method for Android Apps

Majid Hatamian^{1*}, Nurul Momen², Lothar Fritsch², and Kai Rannenber¹

¹ Chair of Mobile Business & Multilateral Security
Goethe University Frankfurt, Germany
{majid.hatamian,kai.rannenber}@m-chair.de

² Department of Mathematics and Computer Science
Karlstad University, Karlstad, Sweden
{nurul.momen,lothar.fritsch}@kau.se

Abstract Smartphone apps have the power to monitor most of people’s private lives. Apps can permeate private spaces, access and map social relationships, monitor whereabouts and chart people’s activities in digital and/or real world. We are therefore interested in how much information a particular app can and intends to retrieve in a smartphone. Privacy-friendliness of smartphone apps is typically measured based on single-source analyses, which in turn, does not provide a comprehensive measurement regarding the actual privacy risks of apps. This paper presents a multi-source method for privacy analysis and data extraction transparency of Android apps. We describe how we generate several data sets derived from privacy policies, app manifestos, user reviews and actual app profiling at run time. To evaluate our method, we present results from a case study carried out on ten popular fitness and exercise apps. Our results revealed interesting differences concerning the potential privacy impact of apps, with some of the apps in the test set violating critical privacy principles. The result of the case study shows large differences that can help make relevant app choices.

Keywords: smartphone apps · case study · security · privacy · Android · privacy policy · reviews · privacy impact · privacy score and ranking · privacy risk · transparency.

1 Introduction

Consumers nowadays frequently use smartphone apps to support and organize various parts of their everyday errands, and accordingly, smartphones have become an indispensable part of our lives. Today’s smartphones are equipped with sensing and recording capabilities such as camera, microphone, fingerprint recognition, proximity sensors, gyroscope, accelerometer, and more. These are embedded into

* This is a post-peer-review, pre-copyedit version of an article published in Lecture Notes in Computer Science book series (LNCS, volume 11498). The final authenticated version is available online at: https://doi.org/10.1007/978-3-030-21752-5_7

the hardware made available to apps and the operating system. As a result, they produce a diverse range of information including sensitive personal information. Importantly, because of their mobile nature and use of wireless communication protocols (e.g. NFC, Bluetooth, 4G, WiFi) to interact with the environment, they are capable to access, use and transmit such sensitive data to remote servers without user interaction or without user insight into what is being transferred. Such a context-sensitive digital ecosystem is highly at risk to produce privacy violations (e.g. unwanted collection, processing, sharing or invasion [17,35]). This makes it quite challenging and difficult for the users to compare apps' privacy aspects and performance and to protect their own privacy. Thus, it is of particular importance to generate transparency by providing quantifiability and thus comparability of apps in regard to their privacy impact [16]

This paper presents a combined method for app privacy analysis and increased transparency that uses several sets of input data. In a joint effort, two research groups [18,19,24] performed a data collection campaign and combined several analysis approaches into the method presented in this paper. We analyze textual privacy policies from app markets. In addition, we extract the use of so-called "dangerous permissions" from the app metadata. We extract and classify end user information on app threats from public app reviews on the Google Play app store. Finally, we monitor app execution by logging app behavior when showing the dangerous permission credentials to the operating system's access control system before they access data sources. The data from these sources then is analyzed and visualized. The method results in tabular and graphical overlays of the input data that can show deviations among privacy policies, reviews, manifestos and actual app behavior. We developed scoring and ranking schemes to compare the level of personal data usage of apps before installation and during installation. To illustrate the method, we show data from a case study with data captured from a set of ten popular fitness apps. Our results enable both *ex-ante* and *ex-post* transparency in the perspective presented in [25], in order to combine the advantages of both concepts, which allows the incorporation of factual app behavior in app choice decisions and app privacy impact evaluation.

Motivation: Which privacy-sensitive data does a mobile app really aim to extract from smartphones? Does the app behavior correlate with the promises of the privacy policy? What are the user's privacy-invasive experiences with the apps? Do the user's concerns reflect correlated privacy threats? And how will a consumer or a public authority decide which app of a set of possible candidates poses the least or an acceptable privacy risk and impact on its users? To answer these questions, we develop a method that extracts data about apps from several sources and prepares the data to enable comparison of app privacy impact.

Contribution: In this paper, we show how data from various sources can be used to assess the potential privacy impact of mobile apps. We further show results from an application of our method to a case study of apps. We identified several privacy issues visible from the data. By providing an understanding of app privacy behavior through data visualization techniques, we show how the data can easily be visualized with each other.

Outline: The rest of this paper is organized as follows: First, we explain our data acquisition and comparison method for privacy impact assessment of mobile apps in Section 2. Section 3 describes our analysis methods used to overlay the data and presents the results of our case study on fitness apps. Then we discuss related work and background relevant for our methodology in Section 4. Finally, we conclude this paper and point out directions for future research in Section 5.

2 Data Acquisition Methodology

Our multilateral privacy impact analysis is based on a four-pillar methodology as shown by Fig. 1. We acquire and process information relevant for app privacy impact from four sources named A1–A4. The sources of information are related to app vendors, end user feedback and actual app behavior measurements. Our method processes both static and dynamic information about each app’s access to personal data. In the following subsections, each pillar is further detailed.

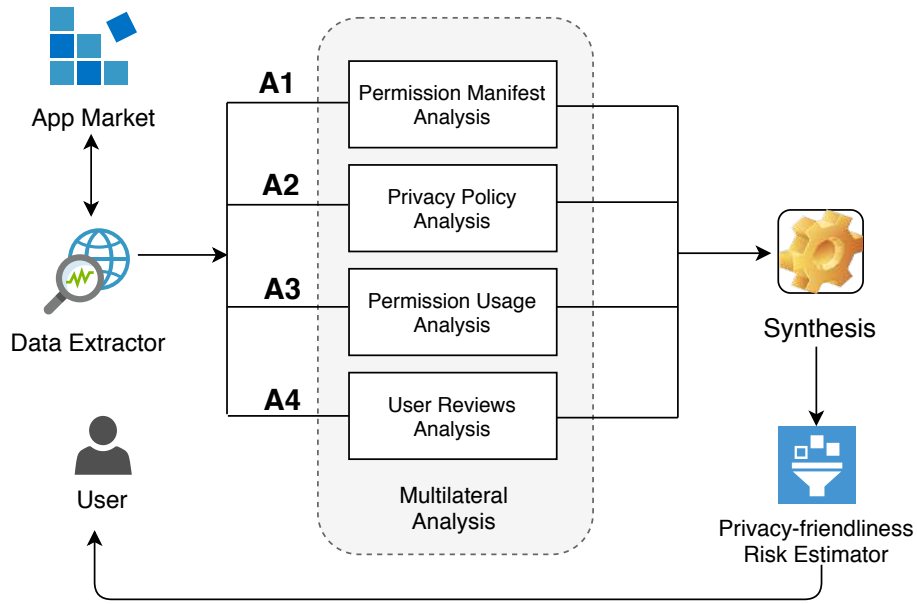


Figure 1: A high-level overview of our multilateral privacy impact analysis approach.

2.1 Permission Manifest Analysis (A1)

We collect app developers’ data access intentions from the apps’ Android manifest. In this app metadata, developers declare use of so-called sensitive permissions

that grant access to data such as call logs, contact lists, sensors or location tracks on smartphones.

Prior to Android 6.0, users had to grant all the requested permissions at install-time and they were not able to revoke those permissions later. Hence, data access was then unlimited for the future. No information about frequency, volume or amount of personal data retrieved and transferred was provided to the data subjects which is still true to some extent for the post-Marshmallow era. However, in Android 6.0 and later versions, Google initiated a new permission manager system where the users are able to revise/revoke permissions at run-time. Although this was an enhancement to give more privacy controls to the users, but still it was not effective. This is mainly because ordinary users mostly do not understand the technical definitions of permission requests [11]. Also, they sometimes value the use of the apps more than their privacy [10,34]. Many apps transfer large amounts of data to remote servers. The access permissions are added by app developers, however the privacy policy prose that should be the base of data subject consent upon installing an app is often very difficult to interpret when looking for cues about what personal data will be extracted from a smartphone. In consequence, it is very difficult to assess the actual consumption of personal data carried out by apps, and thus data subject risk assessment or impact assessment is difficult.

The permissions are usually granted for an app on a permanent basis after initial end user approval upon installation. The user will not learn how often which permissions are being used to access data.

2.2 Privacy Policy Analysis (A2)

As mobile apps are directly dealing with users' personal data, they need to fulfill a certain degree of privacy and security regulation imposed by law e.g. the GDPR [2]. Legislation requires app providers to inform users about their data collection and processing practices in a written privacy policy. Hence, privacy policies are the main source for users to inform themselves about how an app deals with their personal information [32]. In our analysis, we pay attention to privacy policy texts to examine the extent to which they are correlated with what the developer's request (in manifest) and what they do (actual permission usage) in reality. Hence, we also check the extent to which the app privacy policies are actually focusing on the app data collection practices, e.g. whether or not the purpose specification of data extraction based on the dangerous permissions is already clear in the policy text.

2.3 Permission Usage Analysis (A3)

Mobile app users trade their data for service usage in non-transparent ways. Accessibility to user data through permissions gives *carte-blanche*³ access for the

³ Full discretionary power (Merriam-Webster dictionary), Retrieved on November 22, 2018.

app without any constraints. Though the user has the option to revoke granted permissions, the absence of monitoring tools and unexpected consequences such as service exclusion or malfunctions may cause hindrances [5,12,36]. So we measure apps' permission access patterns based on the method described in [23]. We argue that such information can reveal apps' behavior and its impact on individual privacy. It has the potential to assist the user to compare apps based on potential privacy impact and to make decisions based on privacy-friendliness. A comparison matrix or ranking will also be helpful for choosing apps with the least impact for delivering a desired service.

2.4 User Reviews Analysis (A4)

User reviews on app market are an additional source of information regarding app properties. Some contain privacy-related complaints from users. Such complaints can reveal actual privacy risks. Therefore, we try to extract such information from the reviews. However, such information is unstructured and it is quite time consuming to manually code thousands of reviews to gain knowledge about the privacy aspects of apps. Therefore, we exploit this important source of information by automatically collecting reviews and then applying machine learning and natural language processing techniques to extract comments on perceived app privacy problems based on the analysis of user reviews. The resulting data is a mapping of apps into a privacy threat classification. We detect not only a privacy and security relevant user review, but also determine the underlying threat. Based on our already proposed threat catalog [18], we use these threats as the input for the classifier as described in Table 1.

Table 1: Identified threats (shown by T).

#	Threat	Description
T1	Tracking & Spyware	Allows an attacker to access or infer personal data to use it for marketing purposes, such as profiling.
T2	Phishing	An attacker collects user credentials (such as passwords and credit card numbers) by means of fake apps or (SMS, email) messages that seem genuine.
T3	Unintended Data Disclosure	Users are not always aware of app functionalities. Even if they have given explicit consent, users may be unaware that an app collects and publishes personal data.
T4	Targeted Ads	Refers to unwanted ads and push notifications.
T5	Spam	Threat of receiving unsolicited, undesired or illegal messages. Spam is considered an invasion of privacy. The receipt of spam can also be considered a violation of our right to determine for ourselves when, how, and to what extent information about us is used.
T6	General	Comprises all the issues that are not categorized into other threats, such as permission hungry apps, general privacy and security concerns, etc.

3 Multilateral Analysis

This section describes our case study of fitness apps from four different perspectives (A1–A4). It discusses findings and insights gained. Our data collection was performed in October and November 2018. The first two phases of analysis are focused on data sources that are available for ex-ante transparency scenarios. First, apps’ metadata is collected from Google Play store to determine the required permissions that are stated ahead of installation. Second, app’s privacy policy documents are collected and analyzed to adjudicate the cohesion with technical data access intents (manifest data) as described in 2.2. The third and fourth phases are focused on the data sources that are accessible through ex-post transparency scenarios. Ten fitness and exercise monitoring apps (called the app set in the remainder of this text) were chosen based on the top search results on the Google Play app store and were installed and dynamically monitored to measure their permission access requests. Such selection is rationalized as follows: (1) Researchers have raised serious privacy and security concerns resulted from using invasive health and fitness related apps [20,22,28]; (2) Such apps are sometimes underestimated by the users and we intend to highlight the gap between their perception and reality. User reviews of fitness apps can be treated as complementing factor to the technical properties that is measured, which also supports the emphasis on transparency and intervenability by the GDPR [9]. As compared with other popular app categories such as *Lifestyle*, users are not well-aware of the potential negative consequences of using privacy invasive health/fitness-based apps. For instance, in the early 2018, already people were informed about Facebook-Cambridge Analytica data privacy scandal [3]. Hence, it is generally believed that lifestyle-based and social networking apps are the only main potential sources of privacy violations; (3) As a result of extreme proliferation of gadgets and physical activity trackers (such as FitBit), users are currently surrounded by such technologies. Such technological trend is highly dependent on wireless communications between gadgets and smartphones (i.e. health/fitness-based apps) that may potentially impose privacy risks (we can refer to the fitness tracking app gives away location of secret US military bases as a famous example of such dire consequences [4]). The app set is listed on Table 2. Finally, user reviews (collected during the first phase) are analyzed in order to take perceptions and concerns of the user into account. In the following subsections, we explain analysis steps A1–A4 from Fig. 1.

3.1 Step A1: Permission Manifest Analysis

In order to perform tasks in Android, apps can request access to system resources through permissions. The permissions are requested to enable functionality of apps, but they typically exceed this bare-bone minimum requirement, and hence are not privacy friendly. Depending on the resource types, consent from the user is required. There are four types of permissions⁴: *normal*, *dangerous*, *signature*

⁴ <https://developer.android.com/guide/topics/permissions/overview>; [Accessed: 2018-11-27]

and *signatureOrSystem*. *Normal* level permissions allow access to resources that are considered low-risk, and they are granted during installation of any package requesting them. The *dangerous* level permissions are required to access resources that are considered to be high-risk. In this case, the user must grant permission. So-called *signature* level permissions grant access only to packages with the same author. Finally, *signatureOrSystem* level permissions grant both packages with the same author and packages installed in the system receive permission to access specific resources. Every application or app has a *manifest* file which contains information about that particular app (for example - its name, author, icon, and description). It provides information about the required permissions that are requested by the developer. Analyzing the manifest and corresponding permission list offers the primary insight regarding potential personal data access.

Data Collection On the Google Play app store, there is made available public information about apps. Once we obtain the apps' url in the Google Play app store web pages, we can gather the information that we are interested in. We used the scraper in [1] and for each app we retrieved its app ID, title, ratings, number of downloads (installs), app category, permission requests and associated user reviews. Our data set comprises the information of 27,356 apps within *Health & Fitness* category from Google Play. In general, there are 142 distinct types of permissions being extracted across 27,356 apps.

Permission Request Analysis The ten most requested permissions from our app set can be seen in Fig. 2(top). Also, we retrieved the ten most requested permissions corresponding to the ten selected fitness apps (based on the search results, see Fig. 2(bottom)) to examine and compare how different is the requested permissions within the whole category and the chosen app set. As it can be seen, the most and the least widely requested permissions are INTERNET (93.88%) and RECORD_AUDIO (6.55%) respectively. Interestingly, almost all the permission requests (except WAKE_LOCK and VIBRATION) are among dangerous permission requests. When it comes to the chosen set of ten apps, MICROPHONE is substituted by SENSOR. Nevertheless, the rest combination is still intact, however, the percentages and permutations are different.

3.2 Step A2: Privacy Policy Analysis

We also analyzed the declaration of sensitive permission requests by apps to their privacy policy information. For example, we investigated whether or not the app developers claim in their privacy policies that they are going to use a certain sensitive permission. The result is a gap analysis showing the difference between policy declaration and app privileges.

Data Collection We collected the privacy policy texts of the app set. Considering the dangerous sensitive permission request list, two researchers manually




Figure 2: The 10 most requested permissions in: (Top) *Health & Fitness* category, (Bottom) chosen set of popular 10 apps within *Health & Fitness* category.

coded the data and checked the specification of such permissions in privacy policy texts. Due to frequent evolving nature of apps and their corresponding policies, we archived privacy policy documents of apps on 12 November, 2018.

Purpose Specification Analysis Art. 5 (1b) GDPR limits the collection and processing of personal data to “specified, explicit and legitimate purposes” and it says: “*personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes*” [9]. Therefore, it is of particular importance to examine the extent to which the studied mobile apps are fulfilling such requirement. As shown in Table 2, we found 14 incidents where the app developers failed to clarify the need of requesting certain sensitive permissions in their written privacy policy texts (shown by ×).

Table 2: Purpose specification analysis of app privacy policy texts: clarified in the policy: ✓, not clarified: ×, not using that permission: N

App #	CAMERA	SMS	CONTACTS	LOCATION	PHONE	MICROPHONE	SENSOR
Lifesum	×	N	×	N	N	N	×
Endomondo	N	N	✓	✓	✓	N	✓
30dayFitnessChallenge	N	N	×	N	N	N	N
Runkeeper	×	N	×	✓	N	N	N
Pedometer	×	N	✓	✓	×	N	✓
MyFitnessPal	×	N	×	✓	✓	N	✓
Runtastic	N	N	✓	✓	N	×	✓
7minutesWorkout	N	N	N	N	×	N	N
Fitbit	×	×	✓	✓	✓	N	N
Google Fit	N	N	✓	✓	N	N	✓

3.3 Step A3: Permission Usage Analysis

In this section, we present results of a measurement which was conducted in Fall 2018 to determine permission usage patterns of fitness apps in an idle scenario (no user interaction with the app). The app set was installed to observe their activity throughout a period of seven days. In order to do so, apps’ permission access log was collected. Apps accessing lower amount of dangerous permissions are assumed as more privacy-friendly.

Data Collection A prototype probing tool named *Aware* was used for collecting logs of apps’ permission usage [24]. It runs as an Android service and documents apps’ permission access patterns from Android’s AppOpsCommand⁵. Periodically, it checks for the last permission access event by each of the installed apps and writes respective events in a predefined format. Data collection was carried out

⁵ https://android.googlesource.com/platform/frameworks/base/+android-6.0.1_r25/cmds/appops/src/com/android/commands/appops/AppOpsCommand.java; Accessed: 2018-10-23

for one week (starting on 22 October, 2018 and ending on 29 October, 2018). The target apps were installed on a Nokia 5 Android device running on a vendor stock ROM (Android 7.1.1) which was rooted for monitoring. The apps under investigation were not interacted by any user.

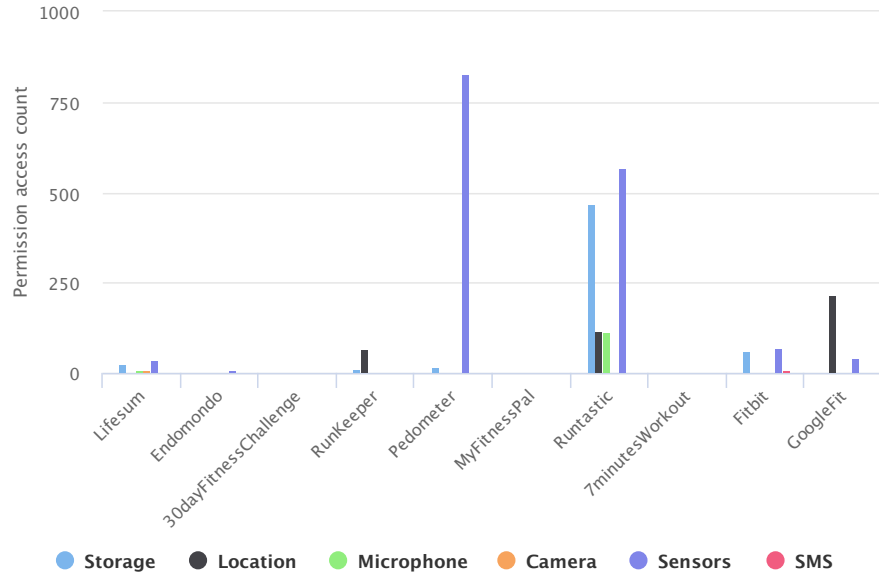


Figure 3: Permission usage: majority of the fitness apps (7 out of 10) kept accessing dangerous permissions, despite having no user interaction.

Fig. 3 shows permission-access activity associated with the unused apps. Accessing to some sensitive permissions such as storage, microphone, SMS and camera while the apps are not being actively used may lead to the following conclusions:

Permission Access Analysis The collected log indicates the intent to access permissions by apps. As idle-time permission access is depicted in Fig. 3, following observations can be drawn from it:

Data Minimization Principle Violation: the permission access events are supposed to be specific to a particular tasks carried out with an app. We found quite the opposite: throughout the experiment period, apps kept accessing permissions. Even though pseudo user installed the apps, their services were not in use. So, resource access by them indicates potential violation of article 5-1(c) of GDPR which states that personal data shall be *adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed* (“data minimization”) [9].

Principle of Least Privilege Violation: principle of least privilege (PoLP) was first proposed as a design principle by Saltzer and Schroeder [33]. According to PoLP, “Every program and every user of the system should operate using the least set of privileges necessary to complete the job.” Clearly, this principle is directly connected to “data minimization” principle, as we observed some apps accessing dangerous permissions which are irrelevant to their intended functionality, for instance in Fig. 3, Lifesum’s usage of CAMERA and MICROPHONE. Also, the need of requesting and accessing such sensitive permissions was unclear in the examined privacy policy texts.

3.4 Step A4: User Reviews Analysis

Crowdsourced user reviews for apps are an additional reference point for identifying privacy threats. It allows us to take the individual’s privacy attitudes into account and map the identified threats to the corresponding cases. We extracted app market user feedback for the app set.

Data Collection Using the tool in [1], we collected a data set consisting of 44,643 user reviews corresponding to the app set from the Google Play app store (in Nov 2018) with a maximum number of 4,500 reviews per app.

Privacy Relevant Complaints Analysis Our goal was to understand what users were posting about privacy issues of apps. We were interested to first extract such information, and then, to determine the granularity of privacy relevant statements (to extract potential privacy threats of apps based on the analysis of their user reviews). Based on our previous work [18], we used the collected data as an input for a trained machine learning algorithm (*Logistic Regression (LR)* implemented in scikit-learn [29]). This ultimately led to a smaller result set. In the end, we detected 1,145 privacy and security-based user reviews. We used recall, precision and F-score metrics to evaluate the performance of the classifier. The values of these metrics show how well the classifier’s results correspond to the annotated results. The observation is that the overall recall, precision and F-score values are of 78.19%, 86.13% and 81.59%, respectively. As the performance analysis of our classification approach is out of the scope of this paper, in the following we mainly focus on the quality of the results (information) that we gathered out of the user reviews. To gain better understanding of the classified user reviews, Table 3 shows some examples regarding the strength of our analysis in distinguishing different types of user reviews and their relevant threat.

In Table 4 and Fig. 4 we report the identified privacy threats associated to each individual fitness app (✓represents the identified threats) and the total number of privacy relevant user reviews per app, respectively. As can be seen, *Runkeeper* and *FitBit* comprise the highest number of threat-related complaints.

Table 3: An example of classified user reviews.

#	Sample user review	T
1	<i>You don't need to spy on my activities outside of this app. they don't care about their customers, they want to ruin the device with horrible bloatware spyware</i>	T1
2	<i>Im still getting warnings that my phone is infected with virus after i update and scan again. If its not going to work why download it. I have very limited memory to use. No need to download stupid apps that dont work</i>	T2
3	<i>SHit!Takes control of device.. why my photo is there??!!</i>	T3
4	<i>Ads are terrible Sorry but the ads are comparing to the website really irritating.</i>	T4
5	<i>Simple interface to use with plenty of features - but pop ups</i>	T5
6	<i>Dangerous! requires unnecessary access to sensitive permissions! Uninstalled</i>	T6

Table 4: List of fitness apps with their respective identified privacy threats (shown by \checkmark).

No.	App name	T1	T2	T3	T4	T5	T6
1	Lifesum	x	x	x	x	x	\checkmark
2	Endomondo	x	x	x	\checkmark	\checkmark	x
3	30dayFitnessChallenge	x	x	x	\checkmark	\checkmark	x
4	Runkeeper	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
5	Pedometer	x	\checkmark	x	\checkmark	x	x
6	MyFitnessPal	x	x	x	\checkmark	x	x
7	Runtastic	\checkmark	x	\checkmark	\checkmark	x	\checkmark
8	7minutesWorkout	x	x	x	\checkmark	\checkmark	x
9	Fitbit	\checkmark	x	\checkmark	\checkmark	\checkmark	\checkmark
10	Google Fit	x	x	x	\checkmark	\checkmark	x

The Most Mentioned Permissions Overall, we found 240 statements corresponding to ten sensitive permissions while some of the privacy relevant user reviews comprise multiple statements referring to a certain permission. Fig. 5 shows the ten user-mentioned permissions out of our analysis concerning the privacy relevant user reviews. The bar chart depicts that the most mentioned permissions are INTERNET, STORAGE and PHONE_STATE (e.g. complaining about access to outgoing calls, phone numbers) being mentioned 46, 44 and 40 times, respectively. In contrast, CALENDAR, CAMERA and MICROPHONE permissions are the least repetitive permissions.

3.5 Synthesis of Analysis

To achieve an overall app privacy impact analysis, we fused the collected data with a scoring algorithm. We presume all permission accesses to be equally risky for privacy. In addition, we treat the different data sets (A1–A4) as contributing equally to privacy impact when fusing the results. In order to do so, total 36 infraction points were set up for calculating cumulative privacy impact score. We assessed the gaps in the privacy policies as defined in Section 2.2. In addition, we monitored idle app data access. Both Table 5 and Fig. 6 show our result—ranking of the app set according to the app privacy impact analysis. We acknowledge that the cumulative sum of privacy impact infraction points lacks some obvious factors e.g. dependability on personal context, subjectivity of risk-perception,




Figure 4: The total number of privacy relevant reviews per selected app.

real-time interaction with apps, individual preferences etc. which remained out of reach for this study due to enormous complexity for adding meaningful weights to impact score and thus, it can be deemed as a limitation.

As the results from four different sources are aggregated into a total privacy impact score as depicted in Table 5, an overall comparison can be drawn from it by ordering from highest to lowest impact score which represents highest to lowest privacy impact. The graphs are presented for each app along with ten dangerous permission groups that could be requested by them (outer blue line in graph). So, an app has the possibility to accumulate total impact score of 36 (10 for requesting permissions, 10 for not clarifying purposes in privacy policy (black segments in graph), 10 for accessing permissions when the app is not in use (red segments in graph) and 6 for identified threats from user review analysis). For instance, in Table 5, *Fitbit*'s privacy impact score is 20 (sum of requested permissions, missing clarifications, usage during idle time and number of identified threats from user review analysis).

From the graphical representation of apps' privacy impact in Fig. 6, it is evident that *30dayFitnessChallenge* and *7minutesWorkout* are more privacy-preserving choices than the rest of the apps. As it is visualized with blue lines (representing permission groups requested in manifest), they are the least permission hungry apps. On the other hand, *Fitbit* is the most permission hungry app (it requests for 9 out of 10 dangerous permission groups). However, apps' privacy policies lack declaration of data processing related to permissions. These discrepancies are visualized with black pie-slices which are placed alongside corresponding permission groups. *Google Fit* and *Endomondo* do not have any discrepancy between their manifest's permission requests and available clarifications in respective privacy policies.

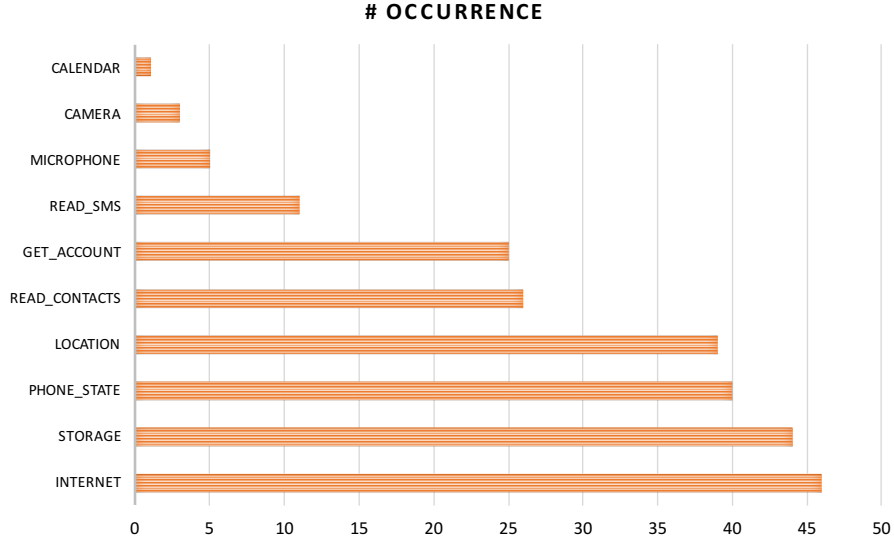


Figure 5: 10 most user-mentioned permissions in user reviews.

Table 5: Synthesis of results from multilateral privacy analysis, ordered by privacy impact score.

App	Privacy impact score (out of 36)	Dangerous Permission Groups Requested (out of 10)	Absent Clarification in Privacy Policy (out of 10)	Idle Permission Usage (out of 10)	Identified Threats from User Reviews (out of 6)
Fitbit	20 (highest)	9	3	3	5 (T1, T3, T4, T5, T6)
Runkeeper	19	6	4	3	6 (T1, T2, T3, T4, T5, T6)
Runtastic	15	6	1	4	4 (T1, T3, T4, T6)
Lifesum	13	5	3	4	1 (T6)
Pedometer	13	6	3	2	2 (T2, T4)
Google Fit	10	5	0	3	2 (T4, T5)
MyFitnessPal	9	6	2	0	1 (T4)
Endomondo	9	5	0	2	2 (T4, T5)
30dayFitness-Challenge	6 (lowest)	2	2	0	2 (T4, T5)
7minutesWorkout	6 (lowest)	2	2	0	2 (T4, T5)



Figure 6: Visual comparison of apps' privacy impact. Headline: App name, privacy impact score out of 36, (parameters in brackets: permissions declared—blue outline, gaps in privacy policy—black segments, data accessed while unused—red transparent segments and T = threat count which is not plotted in graph).

Permission access measured while the app set was installed without user interaction are presented with red areas in the app graphs in Fig. 6. Only three out of ten chosen apps show no idle usage of their listed permissions: `30dayFitnessChallenge`, `7minutesWorkout` and `MyFitnessPal`. The fourth judgment criterion, user review analysis, is not plotted in Fig. 6 due to the fact that it becomes cumbersome for visual representation, but the threat count (T) is considered in total impact score calculation. In Table 5, the identified threats from user review analysis are mapped to the corresponding apps. As it is depicted in the rightmost column for instance, `Runkeeper` is subjected to the most privacy threats that are identified from user reviews, but it ranks second according to the total privacy impact score.

Based on our analysis, an app can be deemed as more privacy-preserving if it requests fewer number of dangerous permissions, has less discrepancy between manifest and available clarification in policy document, has reasonable permission usage during run-time and has fewer threats from user review analysis.

4 Related Work

The assessment of privacy risk and privacy impact suffers from a general shortage of empirical data that provides the basis for privacy risk analysis [13]. Risk calculations are made difficult due to the lack of occurrence and damage information. Analysis therefore looks for other cues, e.g. static properties of program code or code behavior [27,26]. Enck et al. [8] investigated the privacy of smartphone apps by monitoring a set of sensitive permissions, e.g. location, storage, contacts, phone number. In a sample of 311 of the most popular apps downloaded from Google Play, they found five apps that implement dangerous functionalities, and therefore, should be installed with extreme caution. Followed by this study, Enck et al. [7] aimed at understanding of smartphone apps security by proposing a decompiler which recovers Android apps source code directly from its installation image. They analyzed 21 million lines of recovered code from 1,100 free apps using automated tests and manual inspection and it shows the use/misuse of personal/phone identifiers, and deep penetration of advertising and analytics networks. TaintDroid [6] is a method in which the behavior of 30 popular Android apps is studied. The analysis showed that two-third of the apps show suspicious handling of sensitive data and that 15 of them reported users' location to remote advertising servers. *FAIR* [19] is a privacy risk assessment for Android apps and benefits from an app behavior monitoring tool that collects information about accesses to sensitive resources. The authors proposed the calculation of a privacy risk score using a fuzzy logic-based approach that considers type, number and frequency of accesses on resources according to some pre-defined rules. Their analysis on the 15 most popular apps by installation within different app categories on Google Play shows a quantified comparison of apps by reporting to the user the detected privacy invasive events. Although these are important works and provide insights for privacy researchers, but they do not consider the importance of app meta data analysis such as user reviews, privacy policy, manifest declaration, etc.

In [21], the authors investigated the issue of trust when installing a new mobile app. They considered app ratings, reviews and permissions as trust metrics and assessed the trustworthiness of mobile apps. Similar to this, Habib et al. [15] proposed an automatic framework to assess the trustworthiness of mobile apps. Their framework is structured on app’s reputation and state of the art static analysis tools. They evaluated their framework on a data set of some selected apps from the Google Play store that revealed their approach outperforms the existing methods. Neither of these two works studied the privacy-friendliness aspects of mobile apps. Furthermore, they did not investigate the importance of privacy and security analysis of user reviews and they only considered the sentimental aspects of them. Also, the importance of app privacy policy analysis and the correlation between dangerous permission requests (in manifest) and purpose specification (in privacy policy) was not explored. This is why in our work we consider the importance of such aspects and overcome these limitations.

The concept of privacy transparency, in particular ex-post and ex-ante transparency, are presented in detail in [25]. We derived our combined ex-ante and ex-post approach from the ideas discussed in this paper. The privacy impact analysis relates to the principle of multilateral security, which is a security analysis approach that includes all stakeholders’ perspectives and needs in a security analysis [30,31]. The visualization of information is crucial when analyzing and comparing complex information. The data sets in this study are of heterogeneous nature, which poses challenges for visualization. With their systematic overview over visual comparison methods, Gleicher et. al. [14] provided us with useful insights, in particular on overlay encoding of graphs with superposition and explicit encoding.

5 Conclusions and Future Work

In this paper, we presented a method to assess privacy impact of Android apps. The method uses four data sources. We demonstrated the use of the method with a case study performed on ten popular fitness and exercise apps available on the Android app market. Our multilateral methodology allows the assessment and comparison of privacy implication of an app from four different perspectives: a) comparison of apps’ resource requirement, b) assessment of those requirements based on their corresponding privacy policies, c) quantification of their permission access efforts during run-time and d) assessment of privacy concerns raised by users. We combined ex-post and ex-ante transparency perspectives and presented the overlaying results in tabular and graphical overlays as well as in an aggregated privacy impact score which can offer an overview of privacy consequences for a given set of apps. This ranking enables sorting the apps by their potential privacy impact.

The case study found considerable gaps between the privacy policies and the privilege requests and in addition, documented suspicious app behavior of some of the apps in the app set. From this preliminary evidence we conclude that the method has potential in providing transparency about app’s actual intentions

to consume personal data to both end users and regulators. Table 5 and Fig. 6 both show that there are clear differences between app’s access request to data and app vendors’ declaration about their data access intentions. Our results can therefore be used as a base for personal decision-making about continued or future app use.

Our future work will test and refine the method by evaluating the method through studies on app sets for various purposes in diverse contexts. We are also interested to investigate the impact of such visualization and privacy impact analysis on users’ decision making while choosing an app. These steps could include but are not limited to automation of the procedure, prototype development and usability studies. Possibly, our method in the future can support documentation and regulation of privacy violations.

Acknowledgements This research is partially supported by the ALerT project, Research Council of Norway, IKTPLUSS 2017-2021 and by the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 675730 Privacy&Us.

References

1. Google play scraper, <https://github.com/facundoolano/google-play-scraper/>
2. Eu general data protection regulation, <https://eur-lex.europa.eu/legal-content/en/txt/html/?uri=celex:32016r0679>, accessed august 8, 2018 (2016)
3. Facebook data privacy scandal: A cheat sheet, <https://www.techrepublic.com/article/facebook-data-privacy-scandal-a-cheat-sheet/>, accessed jan 11, 2019 (2018)
4. Fitness app strava lights up staff at military bases, <https://www.bbc.com/news/technology-42853072>, accessed feb 01, 2019 (2018)
5. Almuhiemedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F., Agarwal, Y.: Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 787–796. ACM (2015)
6. Enck, W., Gilbert, P., Chun, B., Cox, L.P., Jung, J., McDaniel, P., Sheth, A.N.: Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In: the Proceedings of the the 9th ACM USENIX Conference on Operating Systems Design and Implementation, Vancouver, BC, Canada. pp. 393–407 (2010)
7. Enck, W., Ocateau, D., McDaniel, P., Chaudhuri, S.: A study of android application security. In: the Proceedings of the the 20th USENIX Conference on Security, San Francisco, CA, USA. pp. 21–21 (2011)
8. Enck, W., Ongtang, M., Mcdaniel, P.: On lightweight mobile phone application certification. In: the Proceedings of the the 16th ACM Conference on Computer and Communications Security, Chicago, Illinois, USA. pp. 235–245 (2009)
9. EU Regulation: 679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Off J Eur Union p. L119 (2016)
10. Felt, A.P., Egelman, S., Wagner, D.: I’ve got 99 problems, but vibration ain’t one: A survey of smartphone users’ concerns. In: the Proceedings of the 2nd ACM

- Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM'12), New York, NY, USA. pp. 33–44 (2012)
11. Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E., Wagner, D.: Android permissions: User attention, comprehension, and behavior. In: the Proceedings of the 8th ACM Symposium on Usable Privacy and Security (SOUPS'12), New York, NY, USA. pp. 1–3 (2012)
 12. Franzen, D., Aspinall, D.: PhoneWrap-Injecting the “How Often” into Mobile Apps. In: Proceedings of the 1st International Workshop on Innovations in Mobile Privacy and Security co-located with the International Symposium on Engineering Secure Software and Systems (ESSoS 2016). pp. 11–19. CEUR-WS.org (2016)
 13. Fritsch, L., Abie, H., Regnesentral, N.: Towards a research road map for the management of privacy risks in information systems. In: Gesellschaft für Informatik eV (GI) publishes this series in order to make available to a broad public recent findings in informatics (ie computer science and information systems), to document conferences that are organized in co-operation with GI and to publish the annual GI Award dissertation. p. 1
 14. Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C.D., Roberts, J.C.: Visual comparison for information visualization. *Information Visualization* **10**(4), 289–309 (2011)
 15. Habib, S.M., Alexopoulos, N., Islam, M.M., Heider, J., Marsh, S., Muehlhäuser, M.: Trust4app: Automating trustworthiness assessment of mobile applications. In: 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE). pp. 124–135. IEEE (2018)
 16. Hatamian, M., Serna-Olvera, J.: Beacon alarming: Informed decision-making supporter and privacy risk analyser in smartphone applications. In: To be appeared in the Proceedings of the 35th IEEE International Conference on Consumer Electronics (ICCE), USA (2017)
 17. Hatamian, M., Kitkowska, A., Korunovska, J., Kirrane, S.: “it’s shocking!”: Analysing the impact and reactions to the a3: Android apps behaviour analyser. In: Kerschbaum, F., Paraboschi, S. (eds.) *Data and Applications Security and Privacy XXXII*. pp. 198–215. Springer International Publishing, Cham (2018)
 18. Hatamian, M., Serna, J., Rannenber, K.: Revealing the unrevealed: Mining smartphone users privacy perception on app markets. *Computers & Security* (2019). <https://doi.org/https://doi.org/10.1016/j.cose.2019.02.010>, <http://www.sciencedirect.com/science/article/pii/S0167404818313051>
 19. Hatamian, M., Serna, J., Rannenber, K., Iglar, B.: Fair: Fuzzy alarming index rule for privacy analysis in smartphone apps. In: Lopez, J., Fischer-Hübner, S., Lambrinouidakis, C. (eds.) *Trust, Privacy and Security in Digital Business*. pp. 3–18. Springer International Publishing, Cham (2017)
 20. Hutton, L., Price, B.A., Kelly, R., McCormick, C., Bandara, A.K., Hatzakis, T., Meadows, M., Nuseibeh, B.: Assessing the privacy of mhealth apps for self-tracking: Heuristic evaluation approach. *JMIR Mhealth Uhealth* **6**(10), e185 (Oct 2018). <https://doi.org/10.2196/mhealth.9217>
 21. Kuehnhausen, M., Frost, V.S.: Trusting smartphone apps? to install or not to install, that is the question. In: 2013 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA). pp. 30–37 (Feb 2013). <https://doi.org/10.1109/CogSIMA.2013.6523820>
 22. Martínez-Pérez, B., De La Torre-Díez, I., López-Coronado, M.: Privacy and security in mobile health apps: A review and recommendations. *J. Med. Syst.* **39**(1), 1–8 (Jan 2015)

23. Momen, N.: Towards Measuring Apps' Privacy-Friendliness (licentiate thesis). Ph.D. thesis, Karlstads universitet (2018)
24. Momen, N., Pulls, T., Fritsch, L., Lindskog, S.: How much privilege does an app need? investigating resource usage of android apps. In: 2017 15th Annual Conference on Privacy, Security and Trust (PST). pp. 268–2685. IEEE (2017)
25. Murmann, P., Fischer-Hübner, S.: Tools for achieving usable ex post transparency: A survey. *IEEE Access* **5**, 22965–22991 (2017). <https://doi.org/10.1109/ACCESS.2017.2765539>, <http://ieeexplore.ieee.org/document/8078167/>
26. Paintsil, E., Fritsch, L.: A taxonomy of privacy and security risks contributing factors. In: 6th International Summer School Conference on Privacy and Identity Management for Life, AUG 02-06, 2010, Helsingborg, Sweden. pp. 52–63. Springer (2011)
27. Paintsil, E., Fritsch, L.: Executable model-based risk analysis method for identity management systems: using hierarchical colored petri nets. In: International Conference on Trust, Privacy and Security in Digital Business. pp. 48–61. Springer (2013)
28. Papageorgiou, A., Strigkos, M., Politou, E., Alepis, E., Solanas, A., Patsakis, C.: Security and privacy analysis of mobile health applications: The alarming state of practice. *IEEE Access* **6**, 9390–9403 (2018). <https://doi.org/10.1109/ACCESS.2018.2799522>
29. Pedregosa, F., Varoquaux, G., Gramfort, A., V. Michel, B.T., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, J., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
30. Rannenber, K.: Recent development in information technology security evaluation - the need for evaluation criteria for multilateral security. In: Proceedings of the IFIP TC9/WG9.6 Working Conference on Security and Control of Information Technology in Society on Board M/S Illich and Ashore. pp. 113–128. North-Holland Publishing Co., Amsterdam, The Netherlands, The Netherlands (1994), <http://dl.acm.org/citation.cfm?id=647317.723330>
31. Rannenber, K.: Multilateral security a concept and examples for balanced security. In: Proceedings of the 2000 Workshop on New Security Paradigms. pp. 151–162. NSPW '00, ACM, New York, NY, USA (2000). <https://doi.org/10.1145/366173.366208>, <http://doi.acm.org/10.1145/366173.366208>
32. Reidenberg, J.R., Breaux, T., Carnor, L.F., French, B.: Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkely Technology Law Journal* **30**(1), 39–68 (2015)
33. Saltzer, J.H., Schroeder, M.D.: The protection of information in computer systems. *Proceedings of the IEEE* **63**(9), 1278–1308 (Sept 1975). <https://doi.org/10.1109/PROC.1975.9939>
34. Solove, D.J.: Nothing to Hide: The False Tradeoff between Privacy and Security. Yale University Press (2011)
35. Solove, D.J.: A taxonomy of privacy. *U. Pa. L. Rev.* **154**, 477 (2005)
36. Van Kleek, M., Liccardi, I., Binns, R., Zhao, J., Weitzner, D.J., Shadbolt, N.: Better the devil you know: Exposing the data sharing practices of smartphone apps. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. pp. 5208–5220. ACM (2017)