

Longitudinal Compliance Analysis of Android Applications with Privacy Policies

Saad Sajid Hashmi¹, Nazar Waheed², Gioacchino Tangari¹, Muhammad Ikram¹, and Stephen Smith¹

¹ Macquarie University, Australia
saad.hashmi@hdr.mq.edu.au

{gioacchino.tangari, muhammad.ikram, stephen.smith}@mq.edu.au

² University of Technology Sydney, Australia
nazar.waheed@student.uts.edu.au

Abstract. Contemporary mobile applications (apps) are designed to track, use, and share users’ data, often without their consent, which results in potential privacy and transparency issues. To investigate whether mobile apps have always been (non-)transparent regarding how they collect information about users, we perform a longitudinal analysis of the historical versions of 268 Android apps. These apps comprise 5,240 app releases or versions between 2008 and 2016. We detect inconsistencies between apps’ behaviors and the stated use of data collection in privacy policies to reveal compliance issues. We utilize machine learning techniques to classify the privacy policy text and identify the purported practices that collect and/or share users’ personal information, such as phone numbers and email addresses. We then uncover the data leaks of an app through static and dynamic analysis. Over time, our results show a steady increase in the number of apps’ data collection practices that are undisclosed in the privacy policies. This behavior is particularly troubling since privacy policy is the primary tool for describing the app’s privacy protection practices. We find that newer versions of the apps are likely to be more non-compliant than their preceding versions. The discrepancies between the purported and the actual data practices show that privacy policies are often incoherent with the apps’ behaviors, thus defying the ‘notice and choice’ principle when users install apps.

Keywords: Data privacy, mobile applications, privacy policy, static analysis, dynamic analysis

1 Introduction

“Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others.”³ Between October 2012 and February 2013, Snapchat’s privacy policy said, “We do not ask for, track, or access any location-specific information

³ Alan Westin, *Privacy and Freedom*, 1967.

from your device at any time while you are using the Snapchat application” [1]. However, Snapchat did the opposite by collecting and sharing users’ geo-location information (Wi-Fi and cell-based location data) to its analytics tracking service provider. Federal Trade Commission (FTC) of the United States initiated an investigation, and in December 2014 ordered Snapchat to implement a comprehensive privacy program addressing risks to users’ privacy [2].

In February 2019, FTC fined TikTok 5.7 million USD for illegally collecting children’s personal data [3]. Between 2019-20, in the United States, several federal lawsuits against TikTok were filed citing the harvesting of users’ (including children’s) personal data without consent. In July 2020, these lawsuits were incorporated into a single class-action lawsuit. In February 2021, a settlement was reached where TikTok agreed to pay 92 million USD and stop the collection of users’ bio-metric and location data [4]. Currently, TikTok is facing a privacy lawsuit in the United Kingdom for violating child privacy laws in the collection of personal information and sharing it with third-parties [5].

The examples of Snapchat, TikTok, and various other apps [6] highlight that online service providers frequently do not comply with their privacy policies, despite the presence of regulations related to the disclosure of privacy practices. These regulations include the California Online Privacy Protection Act (CalOPPA) [7], California Consumer Privacy Act (CCPA) [8], and Children’s Online Privacy Protection Act (COPPA) [9] in the US, General Data Protection Regulation (GDPR) in the EU [10], Data Protection Act in the UK [11], and the Privacy Act in Australia [12]. Although there are limitations to privacy policies (such as they are not often read [13], take a long time to read [14], and are hard to comprehend [15]), privacy policies remain the legally recognized standard of protecting privacy based on the “notice and choice” principle [16]. This principle gives users a notice of the data practices performed by the app (or service) before that data is collected. The users then have the choice to *opt-in* (i.e., give permission) or *opt-out* (i.e., decline permission) of giving access to their details.

Several studies have shown the behavior of contemporary apps to be non-compliant with the stated disclosures in their privacy policies [17–23]. Given that the compliance landscape is continuously evolving, we aim to determine whether the non-compliant behavior has changed with time. In this study, we investigate the compliance of Android apps with their privacy policies over time, ranging from 2008 and 2016. We detect and analyze the personally identifiable information (PII) leaked by an app, disregarding the practices made public by the publisher in the app’s privacy policy. We leverage machine learning (ML) techniques to classify the text of apps’ privacy policies and identify the purported practices that collect and/or share data (e.g., phone number). We then uncover the actual data leaks of an app through static analysis (*examining the app code*) and dynamic analysis (*inspecting the network traffic generated by the app*). While relying on existing techniques, this study is, to our knowledge, the first effort to integrate app and privacy policy text analysis to measure how the apps’ privacy conducts have evolved.

The main contributions of this paper are as follows:

- We identify the data collection practices (*privacy practices*) disclosed in 3,151 privacy policies (obtained from 2012 to 2019) of 405⁴ apps using machine learning classifiers. We find that 2,422 privacy policies from 327 different apps disclose at least one of these practices, with an average of 7.86 per policy (§3.1).
- We analyze historical versions of popular apps from 2008 to 2016, based on the union of leaks observed via static and dynamic analysis. The results show an increase in the average number of leaks per app version over time. Surprisingly, the average number of leaks to third-parties rose from 2.7 in 2011 to 4.43 in 2016 (§3.2).
- Our analysis reveals that the compliance of apps with their privacy policy is steadily decreasing from 33.16% in 2011 to 10.76% in 2016. Also, newer versions of the apps are performing more privacy policy violations than their immediately preceding version. For instance, 9.2% of the app versions released in 2016 show an increase in first-party violations compared to their preceding version, whereas only 2.5% show a decrease (§3.3).

The rest of the paper is organized as follows: In Section 2, we describe our data collection approach and present our analysis methodology. In Section 3, we discuss the findings of our analysis. Section 6 reviews the related work, and we conclude and discuss future research directions in Section 7.

2 Dataset and Methodology

Figure 1 depicts our methodology. In the following, we explain the steps involved in our data collection and analysis.

2.1 Dataset

Mobile Apps We select Android apps that are *(i)* popular, *(ii)* have multiple versions, *(iii)* are susceptible to network traffic interception, and *(iv)* have privacy policy in the app’s home-page on Google Play. The popular apps we selected were either in the top 600 free apps based on the Google Play Store ranking, or in the top 50 in each app category, as of 10 Jan 2017. We discard apps with less than four versions and apps for which we can not intercept TLS (Transport Layer Security) traffic (e.g., apps with non-native Android TLS libraries such as Twitter, Dropbox). We also discard apps that do not contain privacy policy on their home-page at Google Play. For the compliance analysis, we collect 5,240 unique APKs that correspond to 268 apps.

We download the selected apps and their previous versions from an unofficial Google Play API [24] that requires the package name and version code to download an APK. We infer the release date of the package from the last modification time of the files inside the APK (such as `AndroidManifest.xml` and `META-INF/MANIFEST.MF`). If the date is incorrect (e.g., before 2008), we refer to third-party services (appcensus.io, apkpure.com) for inferring the release dates.

⁴ see §A for clarification on different number of apps reported.

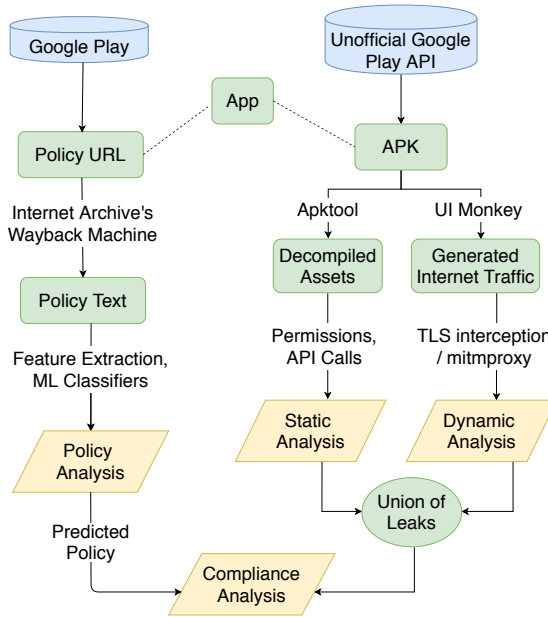


Fig. 1. Overview of our approach. We perform compliance analysis after taking the union of personal identifiable information (PII) leaks detected in both the static and dynamic analysis.

Privacy Policies Google requires Android app developers to disclose the collection and sharing of users’ data [25]. Accordingly, on Google Play Store, each application must provide a link to the privacy policy on its home-page.

For the most updated version of each app, first, we access the URL of the privacy policy from the app’s home-page ([https://play.google.com/store/apps/details?id=\[pname\]](https://play.google.com/store/apps/details?id=[pname]) followed by the package name, i.e., *pname*), and then scrape privacy policy link to extract the privacy policy text.

To obtain the privacy policies of the previous versions of the same app, we leverage the Internet Archive’s Wayback Machine [26], which gives us access to previous snapshots of the app’s privacy policy page. Since 1996, the Wayback Machine has archived full websites, including JavaScript codes, style sheets, and any multimedia resources that it can identify statically from the website’s content. We refer to the single capture of a web-page as a *snapshot*. Wayback Machine mirrors past snapshots of these websites on its servers. Note that we cannot extract the privacy policy content if the URL obtained from the app’s home-page refers to an invalid URL, or if the privacy policy is not found on that URL, or if the privacy policy web-page contains another link that refers to the privacy policy text. Also, the Wayback Machine poses some challenges (such as web-page redirect or varied frequency of archived snapshots) in capturing a spe-

cific web-page [27]. This results in a miss to capture the privacy policy snapshot in that three-month window, and we move to the next three-month window.

There are instances where a single privacy policy is mapped to multiple app versions. This is because privacy policies only became mandatory on Google Play in 2018 [28]. Therefore, a privacy policy residing on a given URL today may not have been residing on the same URL a few years ago. If the Wayback Machine does not archive a policy’s web-page, then the current policy’s web-page (obtained in 2019) is mapped to all the app releases. Furthermore, we are unable to obtain policy web-pages before 2012. Therefore, the app releases prior to 2012 have a high difference in time with the privacy policy date. It may also be possible that the source code of an app is modified when releasing a new version however comes with the old privacy policy.

We use Memento API [29] to capture snapshots at intervals of three months and obtain multiple privacy policies for the same app. Memento API provides the nearest time-stamp for the archived snapshot of a website from the date provided. By comparing the app version’s release date with the snapshot dates of the privacy policy web-page, we obtain the snapshot immediately after the app version’s release date. For example, if an app released versions in Feb’16, Aug’16 and Dec’16, and the app’s privacy policy snapshots are from Jan’16, Jul’16, and Jan’17, then the Feb’16 app version will be assigned the Jul’16 privacy policy. Similarly, Aug’16 and Dec’16 app versions will be assigned the Jan’17 privacy policy snapshot.

2.2 Privacy Policy Analysis

We characterize the identification of privacy practices (i.e., app behaviors that leak PII) from the app’s privacy policy as a supervised classification problem. We initially pre-process the privacy policy text and then apply machine-learning classifiers to predict the privacy practices.

Data Pre-processing: From the privacy policy URL, we obtain the text of the privacy policy by scraping the `< body >` tag of the HTML page. Within the `< body >` tag, we discard unnecessary content nested in other tags (such as `< script >`, `< style >`, `< meta >`, and `< noscript >`). Some archived pages in Wayback Machine have text in a particular format (for example, beginning with “success” and ending with “TIMESTAMPS”) prefixed to the original text. We discard all prefix text prior to the analysis of privacy policy text.

Next, we convert the privacy policy text into segments (or paragraphs) by ensuring that a single paragraph does not contain less than fifty characters, except if it is the last paragraph of the privacy policy. Previous work has shown that practices are better captured from the policy text if the text is broken down into segments [30, 31]. If a text segment (other than the last segment) has less than fifty characters, it is likely to be the heading of the next paragraph; therefore, we merge it with the next text segment. We also combine two adjacent text segments if their combined length is less than 250 characters. After converting the text into segments, we lowercase all segments, and normalize whitespaces and apostrophes (for example, we change words like *haven’t* and *don’t* to *have not*

and *do not* respectively). We finally discard non-alpha and non-ASCII characters and also remove single-character words from the text segments.

Finally, we extract vectors for the segment text by taking the union of TF-IDF [32] vector and a vector of hand-crafted features obtained from [33]. We choose TF-IDF because it is the most popular weighting scheme in the domain of text mining and information retrieval [34, 35]. We create the TF-IDF vector using `TfidfVectorizer` [36] with English stop words from the Natural Language Toolkit (NLTK) corpus. The vector of hand-crafted features comprises boolean values corresponding to the presence or absence of custom strings in the training dataset. For example, we use the strings `cookie`, `web beacon`, and `tracking pixel` to indicate the disclosure of `Identifier.Cookie` practice.

Training Data: To build a training dataset, we leverage the work of Zimmeck et al. [22] where they created a privacy policy corpus of 350 most popular mobile apps (APP-350 corpus [37]). All the 350 apps selected have more than 5 million downloads. Legal experts annotated the privacy policies to identify the privacy practices mentioned in the policy text. The annotated label on a privacy policy comprises three parts (or tiers): *(i)* a general or specific practice (twenty-eight unique practices in total), *(ii)* whether that practice has been performed or not, and *(iii)* whether that practice is associated with first or third-party. In first-party practice, the PII is accessed by the code of the app itself, and in third-party practice, the PII is accessed by third-party libraries (such as analytics, advertisements, or social networks). The annotation labels are assigned to the segments (or paragraphs) rather than the whole policy text. The union of labels for the segments provides the disclosure in the whole policy text. We randomly split the APP-350 corpus into training data (n=250) and test data (n=100).

Classification Problem: For a given privacy policy, we aim to detect the disclosure of privacy practices. A privacy practice comprises three components: *(i)* type of PII (e.g., `Contact`), *(ii)* negation or approval of the practice being performed (`Performed` or `Not Performed`), and *(iii)* party (`1stParty` or `3rdParty`). We characterize this task as a multi-label text classification problem.

We sub-divide the classification task to identify: *(i)* PII type, *(ii)* procedure i.e., practice has been performed or not, and *(iii)* first-party or third-party. If we make classifier for each unique combination, we will require 112 classifiers for each combination of PII type, perform/not perform, and first/third-party. The limitation of this approach is that the training samples for most of these combinations are limited (less than 100). Therefore, after sub-dividing the classification problem, we only require thirty-two classifiers (twenty-eight for unique PII types, two for procedure (`Performed` or `Not Performed`), and two for *parties*). For example, `Contact_Email_Address Performed 3rdParty` will be assigned to a text segment for which `Contact_Email_Address`, `Performed`, and `3rdParty` classifiers all return a positive value for that text segment. We consider a policy text discloses a privacy practice if at least one segment in the text returns positive values for at least one PII type classifier, `Performed` classifier, and `1stParty` and/or `3rdParty` classifiers.

We utilize One-vs-the-rest (OvR) [38] classification strategy and test our approach with the Multinomial Naive Bayes [39], Logistic Regression [40], and Linear Support Vector Classifiers (SVC) [41]. We select Linear SVC as our classifier for unseen policies’ text due to its superior performance on the test dataset (n=100). In particular, the average F1 score (%) for all classifiers with Multinomial Naive Bayes and Logistic Regression is 17.37% and 67.54%, respectively, while for Linear SVC, we obtain an average F1 score of 73.75%. Table 7 shows the performance of our classifiers. Our approach achieves high accuracy, ranging from 91.16% to 100% for all the classifiers.

2.3 Static Analysis

Resources of an app that are required to run it on a device are bundled together in an Android Package Kit (APK). In the static analysis of mobile apps, after downloading the APK, we decompile it using Apktool [42]. Decompiling APK yields the assets (including byte-code in the DEX format) and metadata (in the XML format). After disassembling .dex files into smali format, we extract the API calls. We also extract the app permissions from the `AndroidManifest.xml` file. If an app needs access to a resource outside its sandbox, then it will request the appropriate permission in the `AndroidManifest.xml` file. API calls that do not have the required permissions are not executed, for example, the API `android.telephony.TelephonyManager.getImei` requires `READ_PRIVILEGED_PHONE_STATE` permission to retrieve the IMEI (International Mobile Equipment Identity). In particular, we utilize the Android APIs from [22] and obtain their required permissions from Android API reference [43]. If we observe an API call for a particular privacy practice in the source code, we check if the required permissions are also requested in the app manifest file. We identify that practice is being performed (for a given privacy practice) if both the relevant API call(s) and the required permission(s) exist.

We distinguish the API *function* calls into system calls, first-party calls, and third-party calls. Calls made by the Android class are classified as system calls. We differentiate between first and third-party calls by comparing the classes based on Java’s reversed internet domain naming convention [44]. The `.smali` file’s package name is matched with the app’s package name. If (i) both top and second-level domains match, or (ii) the `.smali` file’s package looks obfuscated (e.g., `b/a/y.smali`), we classify the API call as first-party. Otherwise, it is classified as a third-party call. In our analysis, we only consider the API calls made by first or third-parties.

Besides obtaining the third-party “domains” invoking the API calls, we need to extract the corresponding company names in order to check if these appear in the privacy policy text or not. Since a company can acquire multiple domains, we leverage the previous work of Binns et al. [45] in the domain-company ownership. For example, if the API call is from `adsense.com` then we check for the existence of terms “adsense”, “google” and “alphabet” in the privacy policy text. If we find any of the terms in the policy text, we consider the third-party domain as disclosed.

2.4 Dynamic Analysis

A limitation of static analysis is its inability to capture dynamically loaded code and analyze obfuscated function calls [46]. To overcome this limitation, following the work by Ren et al. [47], we complement our approach with the dynamic analysis of apps. We employ ReCon [48], a transparency control tool that relies on the machine-learning classification to identify leaks of privacy-related information in the mobile-app traffic [49]. In particular, using a public dataset of annotated traffic flows [50], we train a machine-learning classifier (C4.5 decision tree) to predict whether traffic flow is leaking PII. As a feature set, the classifier takes a concatenation of the text of URI, Referrer, postData text, and all other HTTP headers in the flow. The validation accuracy of the classifier is 97.2% (AUC=0.987), with 97.8% precision and 96.6% recall. For those flows predicted to leak (any) PII, we extract the performed privacy practices by matching the feature set against a predefined set of keywords and regular expressions from ReCon [51].

The gold standard for identifying privacy leaks is by manually logging into the apps and interacting with them. However, this approach is impractical at scale. To automate the analysis, we rely on Android’s UI/Application Exerciser **Monkey** [52], a command-line tool that generates pseudo-random user events such as swipes, clicks, or touches. While running an app, we use **mitmproxy** [53] (a TLS-capable interception proxy) to capture all app-generated traffic (HTTP and HTTPS) on a dual-stack [54] WiFi in our testbed. Prior work showed that synthetic usage patterns could lead to underestimating the number of privacy leaks compared to manual (human) interactions [49] since random streams of Monkey events may miss some function calls. While this is a common drawback of automation approaches, Android’s Monkey exhibits the best code coverage among existing automation tools [55].

In our effort to capture most—if not all—the privacy leaks produced by an app, we take the union of leaks from static and dynamic analysis. Since labels are different in the two cases, we map dynamic leak labels onto the static ones. To this end, we compare the descriptions of static leak labels (*see* Table 6) with the ones of dynamic leak labels in [47]. Table 1 lists the conversion between the leak labels. We keep the static leak labels as our final ones since they match the labels in the privacy policy text, making them convenient for the following subsection.

2.5 Compliance Analysis

In the compliance analysis, we compare the leak labels with the privacy policy labels. We consider a compliance violation to have occurred if there is a positive value for a leak label and non-positive value(s) for the corresponding policy text label(s). For some dynamic leaks labels in Table 1, a positive value for a leak label is mapped to multiple labels in static leaks. For example, a positive value for **gsf id** will result in mapping of positive values for both **Identifier Ad ID** and **Identifier Cookie**. Therefore, if we compare individual combined leak labels only, then we may come across unintended violations. For example, if there are positive values for **Identifier Ad ID** and **Identifier Cookie** in the leak

Table 1. Label mapping between PII leaks from static analysis (*static* leaks) and leaks from dynamic analysis (*dynamic* leaks).

Static Leaks	Dynamic Leaks
Contact	firstName, lastName
Contact E Mail Address	email (+ hash)
Contact Password	password (+ hash)
Contact Phone Number	phone number
Demographics Gender	gender
Identifier	hardware serial (+ hash)
Identifier {Ad ID, Cookie}	gsf id, advertiser id (+ hash)
Identifier Device ID	android id (+ hash)
Identifier IMEI	imei (+ hash)
Identifier {IMSI, Mobile Carrier, SIM Serial}	sim id
Identifier MAC	mac addr (+ hash)
Location {Bluetooth, Cell Tower, GPS, IP Address, WiFi}	location

labels, but only a positive value for `Identifier Cookie` in the policy text labels, then our system would be flagging `Identifier Ad ID` as a violation. Thus, to avoid these unintended violations caused by our mapping table, we consider a violation to have occurred if the policy labels do not return a positive value for any of the similar labels (both `Identifier Ad ID` and `Identifier Cookie` in this example).

For APKs that leak PII to third-parties, we also check for the existence of those third-party domains in the APK’s privacy policy text. We examine the domains that are not frequently mentioned in the privacy policies. Note that the mention of the domain in privacy policy does not necessarily imply that the policy text is indicating the sharing of PII with that domain, but the absence of the domain from the policy text definitely implies that the company or domain of the organization with which the app shares PII is not disclosed. Furthermore, by distinguishing apps according to their categories, we identify the most transparent categories (complying with the privacy policies) and categories that have the most compliance issues.

3 Analysis and Results

In this section, we discuss our analysis of apps’ privacy policies and then present leakages observed via our methodology. Finally, we present an analysis of apps compliance with their privacy policies.

3.1 Privacy Policy Analysis

We obtain 3,151 unique privacy policies (P.P.) from the Internet Archive and map them to 7,998 different app versions of 405 apps. We term an app’s privacy policy as unique if it is obtained on a different date for an app. For example, if we collect two privacy policies of an app from Internet Archive on two separate

dates, we term that two P.P. as unique, even if the contents of both the policies are identical.

To map the P.P. to the different app versions, we bind each APK to the app’s privacy policy immediately after the APK release date. We note that 42.6% of the analyzed APKs have a time difference of fewer than 12 months between the APK release date and the P.P. date (*see* §B).

The collected privacy policies contain 143,783 policy segments (or paragraphs) in total. We restrict this set by considering only those segments that return positive values on (i) at least one of the twenty-eight practice classifiers (e.g., Contact), (ii) procedure (i.e., the practice described in the segment is being performed or not performed), and (iii) parties (i.e., 1st Party or 3rd party). Out of a total of 143,783 policy segments, 19,371 segments return a positive value for the above three categories. We term these segments as *valid segments*. The valid segments span across 2,455 privacy policies mapped to 329 apps. We will rely on this subset (valid segments) to measure privacy policy violations in the remainder of the analysis.

Figure 2a shows the distribution of the valid policy segments by the procedure. A vast majority (97.48%) of valid policy segments have a positive value for the **Performed** classifier. These segments (n=18,882) span across 2,422 P.P. mapped to 327 apps. This also includes the segments that return positive values for both **Performed** and **Not Performed** classifiers. A segment is not contradictory if it has positive values for both **Performed** and **Not Performed**. For example, the segment “*we will only store your email and will not collect your location*” can have positive values returned by both the classifiers. In such cases however, our classifiers are unable to distinguish between **Performed** PII type(s) and **Not Performed** type(s). Therefore, to overcome this limitation and prevent falsely flagging a policy for a violation, we assign all flagged PII types as **Performed**. Similarly, for the segments that return positive values for both the parties (n=2,944), we consider the flagged PII types(s) to be **Performed** by both the parties. Among the segments that return positive value for the **Performed** classifier (n=18,882), most have at least one practice being performed by the first-party (88.14%).

The segments that return a positive value for the **Not Performed** classifier only (n=489) span across the P.P. of 89 apps and 382 APKs. Among them, 33 APKs comprising eight apps have segments that only return a positive value for the **Not Performed** classifier, i.e., these APKs’ P.P. only state that PII-revealing practices are not performed. The three most common practices disclosed in the P.P. are **Identifier Cookie**, **Identifier IP Address**, and **Contact E Mail Address**. We empirically observe that the number of practices disclosed in the analyzed P.P. is 19,038. This number implies an average of 7.86 per APK, or 7.75 if we also include the APKs whose policy segments only return a positive value for the **Not Performed** classifier.

For a given year, a P.P. of an app can have multiple versions (based on the date the P.P. was obtained). Among all the P.P. that disclose the collection of PII to 1st party, 86.5% comprise multiple versions in the same year (86.1% for 3rd

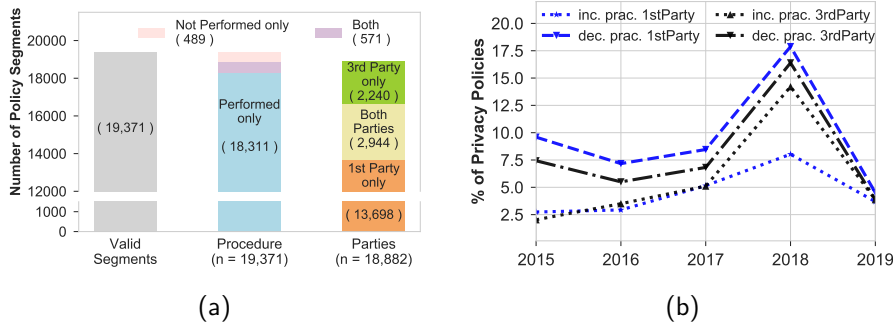


Fig. 2. (a) Breakdown of valid segments ($n=19,371$) into Procedure (Performed / Not Performed). Valid segments for which Performed classifier returns a positive value ($n=18,882$; Performed + Both) are broken down further into parties (1st / 3rd). (b) Percent of P.P. versions in the given year showing an increase/decrease in the number of practices disclosed. Among the newer versions of P.P. each year, more number of them continue to disclose less PII.

party, respectively). Among these multiple versions, we compare the P.P. of an app to their preceding version to identify if they are disclosing the collection of more or less PII. Figure 2b shows the annual disclosure trend of P.P. obtained in the years 2015-2019. We do not consider P.P. before 2015 due to fewer versions available (< 10 annually). We observe the prevalence of P.P. with fewer PII disclosures compared to their preceding version. In particular, 17.9% of P.P. in 2018 show decreasing numbers of 3rd party disclosures (respectively, 16.4% 1st party disclosures), whereas only 8% of P.P. are found with more 3rd party disclosures (respectively, 14.2% 1st party disclosures). For apps whose behavior (leaks) remain uniform, the drop in the number of disclosures in P.P. can have potential compliance issues.

3.2 Analyzing Leakages

Static Analysis: We analyze 7,741 different APKs (or app versions) of 350 different apps, from which we extract the PII-sensitive API calls. Among these calls, 24.2% ($n = 28,001$) are made by the first-party whereas 75.8% calls ($n=87,497$) are initiated by third-party services or domains.

Figure 3a shows the distribution of these API calls to different types of PII. `Identifier Device ID` and `Identifier Mobile Carrier` are the top two PII that API calls (from both first and third-parties) are collecting.

Contrasting Static and Dynamic Leaks: Table 2 shows the annual distribution of apps and versions (APKs) that report PII leak(s) through static analysis, dynamic analysis, and both.

We observe significant differences in the sets of PII-leaks found by static analysis and dynamic analysis, with the former revealing more leaks than the latter in the majority of cases ($>85\%$ of analyzed APKs) as shown in Figure 3b. For example, `com.ace.cleaner` version 6.0 leaked `Identifier Device ID`, `Contact E Mail Address`, `Identifier MAC`, and `Identifier Mobile`

Table 2. Number of unique apps (and versions i.e., APKs) with at least one PII leak. We report the leaks revealed by static and dynamic analysis. We also enumerate the apps (and APKs), which are common to both the analysis. We only consider the apps with retrieved privacy policy.

Year	(A) Static		(B) Dynamic		(A \cap B)	
	#Apps	#APKs	#Apps	#APKs	#Apps	#APKs
2008	3	33	5	44	2	30
2009	6	32	6	18	3	4
2010	24	102	17	34	13	20
2011	47	253	39	119	26	69
2012	87	550	66	197	53	171
2013	136	868	81	334	66	299
2014	173	1302	110	525	95	476
2015	272	2035	208	1018	178	847
2016	301	2485	235	1492	197	1242




Fig. 3. (a) Number of sensitive API calls observed in the static analysis. (b) Difference in number of static leaks and dynamic leaks. At least 85% of analyzed APKs have more leaks discovered in the static analysis than in the dynamic analysis.

Table 3. Summary of PII Leaks from analyzed apps. In columns 2, 3, and 6, the numbers in parenthesis i.e., () represent the number of versions of apps (APKs).

Year	#Apps (#APKs)	Leaks to 1st Party			Leaks to 3rd Party		
		# Apps (#APKs)	#Leaks	#Leaks per APK	#Apps (#APKs)	#Leaks	#Leaks per APK
2008	6 (47)	4 (34)	116	3.41	5 (34)	0	0
2009	9 (46)	6 (29)	116	4.00	6 (21)	12	0.57
2010	28 (116)	20 (84)	252	3.00	21 (63)	187	2.97
2011	59 (303)	39 (169)	558	3.30	53 (269)	726	2.70
2012	99 (576)	67 (354)	1127	3.18	94 (525)	1531	2.92
2013	149 (903)	99 (637)	2217	3.48	141 (835)	3202	3.83
2014	188 (1351)	132 (940)	3289	3.50	182 (1306)	5780	4.43
2015	302 (2206)	207 (1485)	5475	3.69	300 (2155)	9228	4.28
2016	339 (2735)	234 (1841)	7184	3.90	336 (2707)	11998	4.43




Fig. 4. Cumulative Distribution Functions (CDFs) of leaks for (a) 3rd party. Annual CDF of violations for (b) 1st party, and (c) 3rd party. Over time, the total number of leaks from APKs and PII violations are increasing.

`Carrier` to third-parties in the static analysis whereas only `advertiser id`, and `location` were leaked to third-parties in the dynamic analysis. Part of the PII leaks is missed by dynamic analysis due to limitations of Android’s UI/Application Exerciser Monkey in triggering all PII-related API calls. At the same time, we observe that static analysis could not detect some PII leaks (user’s ad ID, cookie ID, location data sharing) as these appeared in dynamically loaded code or obfuscated function calls. The limitations of both static and dynamic analysis demonstrate that only the union set of leaks from the static and dynamic analysis can allow for effective or useful profiling of the privacy-related behavior of an app.

Besides, we examine the prevalence of detection from either static or dynamic analysis for different PII leaks in the period 2012-2016 – we do not analyze the years before 2012 due to insufficient data (<350 APKs). We empirically observe that the most commonly accessed PII is the `Identifier Device ID` by first-parties and `Identifier Mobile Carrier` and `Identifier Device ID` by third-parties.

Trend of Leaks: We then take the union of PII leaks observed in static and dynamic leaks using the mapping table (*see* Table 1). The summary of combined leaks, shown in Table 3, suggest that the average number of PII leaks per APK has risen consistently since 2012, from an average of 3.18 to first-parties (respectively, 2.92 to third-parties) in 2012 to 3.9 (respectively, 4.43 to third-parties) in 2016. Another interesting takeaway is the rise of leaks to third-parties in comparison to the first-parties. Till 2012, apps leaked PII mostly to the first-parties, and since then, they have been leaking mostly to third-parties. We can also observe this trend in Figure 4, which shows a more significant increase over time for the third-party leaks (Figure 4a) compared to the case of first-party leaks. For instance, 62.9% of the APKs exhibit less than three leaks to third-parties in 2012, compared to 27.3% in 2016. Similarly, 21.7% of the APKs exhibit more than five leaks to third-parties in 2012, which increased to 31.1% in 2016. This trend can be attributed to the increasing embedding of ad and tracking libraries in the apps [56, 57].

Table 4. Summary of contradictions or violations of analyzed apps with their privacy policies. In column 2, the numbers in parenthesis i.e., () represent the number of app versions (APKs).

Year	#Apps (#APKs) with ≥ 1 valid seg.	#(%)APKs comply. with P.P.	#Leaks Total	#Leaks viol. P.P.	#Leaks per APK viol. P.P.	#Leaks to 1st P. viol. P.P.	#Leaks to 3rd P. viol. P.P.
2008	3 (11)	11 (100)	8	0	0	0	0
2009	5 (24)	18 (75)	57	9	0.38	0.38	0
2010	18 (80)	37 (46.25)	313	71	0.89	0.56	0.32
2011	37 (187)	62 (33.16)	820	240	1.28	0.70	0.59
2012	69 (364)	117 (32.14)	1622	556	1.53	0.65	0.88
2013	96 (584)	104 (17.81)	3473	1264	2.16	0.77	1.39
2014	126 (860)	79 (9.19)	5332	2160	2.51	1.00	1.51
2015	195 (1381)	149 (10.79)	8640	3909	2.83	1.10	1.73
2016	225 (1748)	188 (10.76)	11338	5210	2.98	1.04	1.94

3.3 Compliance Analysis

Disclosure of practices performed: In the compliance analysis, we compare the combined (static and dynamic) leaks of an APK with the practices reported in the privacy policy. A privacy policy violation occurs if a leak observed in static/dynamic analysis is not mentioned in the privacy policy. For example, our analysis finds that the app `com.fitbit.FitbitMobile` [58] with the version code 2182996 (APK release date: 22 Jun 2016) has privacy policy (dated: 16 Sep 2016) violations for `Identifier Device ID`, and `Identifier Mobile/Sim` practices being performed by the third-parties. On manually inspecting the privacy policy text, we do not find the disclosure of the practices mentioned above for third-parties. For disclosing data to third-parties, the app’s policy states: “*Data That Could Identify You Personally Identifiable Information (PII) is data that includes a personal identifier like your name, email or address, or data that could reasonably be linked back to you. We will only share this data under extremely limited circumstances*”. This generic clause may seem to be allowing the app developers to share the user’s data, but Android’s policy on user data states that the developers must describe the data being collected and explain their usage [25]. This means that the PII sharing from the Fitbit app does not comply with the declared public privacy policy of the app.

Table 4 lists the annual summary of violations observed in APKs. These violations only refer to those APKs whose privacy policy contains at least one valid segment i.e., a segment that returns positive values on (i) at least one of the twenty-eight practice classifiers (e.g., Identifier Cookie), (ii) procedure (i.e., the practice being described in the segment is being performed or not performed), and (iii) parties (i.e., 1st Party or 3rd party). Most of the policy text segments (124,412/143,783) that we classified did not return positive values for the above three categories (see §3.1). For these cases, our machine-learning classifier could not determine whether the policy text presents a practice as “performed” or “not performed” by the parties. From Table 4, we can observe that the compliance




Fig. 5. Percent of APKs released in the given year showing an increase/decrease in the number of violations compared to their preceding version (APK). Newer versions of apps continue to leak more PII that violate privacy policy.

of APKs with their P.P. has decreased considerably from 33.16% in 2011 to 10.76% in 2016. The average number of leaks per APK that is not disclosed in the P.P. is also steadily increasing. We can observe this trend in Figure 4b and 4c, showing that the number of P.P. violations is consistently increasing overtime for a significant fraction of APKs.

We empirically note that `Identifier Mobile/Sim` (grouping of `Identifier IMSI`, `Identifier Mobile Carrier`, and `Identifier SIM Serial` based on the mappings in Table 1) annually contributes to more than 25% of first-party and more than 30% of third-party violations. `Identifier Device ID` is next, comprising more than 20% of annual violations. This suggests that privacy policy violations are often due to apps *fingerprinting* users' mobile devices without revealing it.

Compliance of an app across newer versions: For a newer version (APK) of an app released, it can have (i) equal number, (ii) higher number, or (iii) smaller number of privacy policy violations compared to the preceding version of the given app. Figure 5 shows the annual compliance trend of APKs released in the years 2012-16 compared to their preceding version. For example, if an app released versions in Oct'15, Dec'15, Feb'16, Jul'16, and Nov'16, then for the year 2016, the number of violations of APK released in Nov'16 will be compared with Jul'16, Jul'16 will be compared with Feb'16, and Feb'16 will be compared with Dec'15. For the set of app versions released in 2012-2016, we observe the prevalence of apps with more privacy policy violations compared to their preceding version. In particular, 9.1% of APKs from 2016 show increasing numbers of third-party violations (respectively, 9.2% first-party violations), whereas only 3.8% of APKs are found with fewer third-party violations (respectively, 2.5% first-party violations).

Disclosure of 3rd Party Domains: Given the set of PII leaks to third-party domains, we verify whether those domains are mentioned in the P.P. text. Table 5 shows the distribution of P.P. based on the disclosure of third-party domains. We note that a vast majority of P.P. (82.2%) do not mention at least one domain with which they share PII.

Table 5. Breakdown of privacy policies based on the disclosure of 3rd party domains with which their APKs share PII.

% (#) of APKs with Privacy Policies (P.P.) mentioning		
ALL 3rd party domains	NONE 3rd party domains	PARTIAL 3rd party domains
17.8% (1,419)	23.85% (1,901)	58.35% (4,651)

Figure 6a shows the annual distribution of the number of domains that are not mentioned in the P.P. We observe a significant increase in the number of third-party domains that app publishers do not disclose despite being involved in PII sharing. In particular, while in 2012, almost 70% of the released APKs had no more than one 3rd party domain “missed” by their privacy policies; in 2016, this percentage dropped to less than 30%. Figure 6b shows the twenty most recurrent domains that are not disclosed in the privacy policies. Overall, these twenty domains account for 49.5% of the total instances where third-party domains are not disclosed. Each of these domains provides a library for analytical services, advertisements, social networking, or utility/developer tools to the app. The most frequently not mentioned domain in the privacy policies by app publishers is `facebook.com`, accounting for more than a quarter (27.8%) of the analyzed APKs. It is followed by `flurry.com` (19.7%) and `google.com` (18.9%).

4 Discussion

The increase in the number of violations over time is concerning. The causes of increased violations can be that (i) the P.P. has shrunk, i.e., newer versions of P.P. of an app are disclosing the collection (or sharing) of less number of PII, (ii) the behavior of app has changed, i.e., newer versions of an app are leaking more PII, (iii) or both. We observe in Figure 2b that among the newer versions of

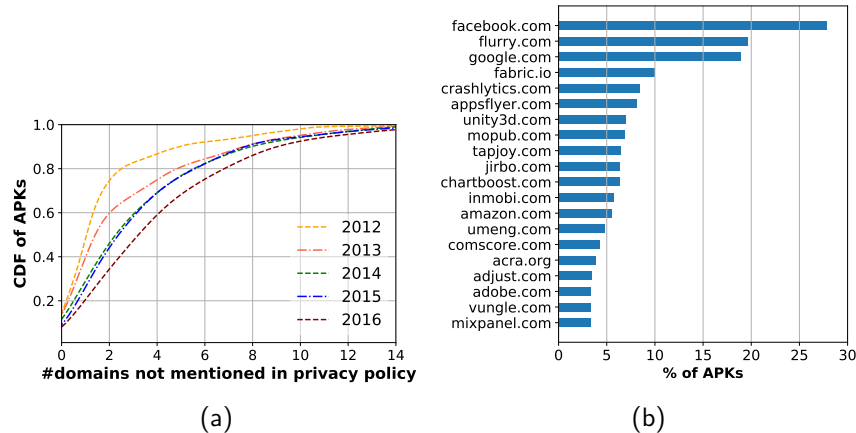


Fig. 6. (a) CDFs of number of domains not mentioned in the privacy policy of APKs. Each subsequent year, APKs share PII to more 3rd party domains that are not disclosed in their privacy policy. (b) Prevalence of undisclosed domains among the privacy policies of APKs.

P.P., the number of P.P. that had a decrease in the number of disclosed practices compared to the preceding version were in greater proportion than those that had an increased disclosure of PII. Furthermore, as evident from Table 3, the number of leaks has also been on the rise, particularly to third-party domains. With the rise of third-party advertisements and analytics services, this comes as no surprise. The increase in the concealment of third-party domains (most recurring domains being advertisement and analytics) from the P.P. text (*see* Figure 6a) lends further support to this argument.

5 Limitations

Despite our best efforts to longitudinally analyze apps’ compliance, there are still some limitations to our work.

To scale up the study and cope with a large number of apps/versions, we leveraged automated analysis tools and machine learning classification. While this may result in “false positive” detection of privacy violations, the risk of false positives is, in fact, limited: static app analysis techniques are deterministic (they rely on pattern matching), and dynamic privacy leak detection has high validation accuracy (above 97%).

The use of machine learning and natural language processing for analyzing the privacy policy texts has some limitations. Privacy policies with generic statements or subjectively redundant words are arduous to interpret without human intervention. The use of subjective verbiage in handling users’ data is also not recommended by Android [25].

In this study, the level of severity of privacy leaks is a fixed assumption. For example, some users may find the location a more intrusive leak than email id and vice-versa. Furthermore, the gap between the app version release date and the P.P. date (policy snapshot that matches the release date) can be several months. This gap is because the Wayback Machine could not find a P.P. snapshot closer to the release date. Therefore, it may be possible that the policy under consideration is significantly different from the irretrievable one and thus may not be accurately reflecting the app version’s data practices.

6 Related Work

In recent years, there have been many research studies on the privacy implications of web tracking, and privacy protection [59–63]. In privacy policies, apps’ developers must declare all the permissions an app will require to perform the tasks; however, previous study has shown that websites rarely disclose “Do Not Track” DNT in their privacy policies, and most websites do not comply with DNT even after disclosing it [64].

Zimmeck et al. [22], Slavin et al. [17], and Wang et al. [18] employed static analysis for analyzing privacy policy violations in Android apps. Zimmeck et al. [22] and Slavin et al. [17] used Android API calls to identify the collection of users’ PII from mobile devices. Using native code, Wang et al. [18] check data transparency for data provided by users via GUI components. Han et al. [21]

used static and dynamic analysis to compare free and paid versions of the apps for data leaks, embedded third-party libraries, and sensitive permissions. Reyes et al. [19], Okoyomon et al. [20], and Andow et al. [23] leveraged dynamic analysis for identifying contradictions in the behavior of apps with the privacy policies or regulations.

Perhaps, a much closer work to our’s is by Ren et al. [47], where they monitored the network traffic to detect PII leaks of 512 Android apps across different versions. They show that apps leak PII to more third-party domains over time. We build on this work by conducting a compliance analysis of apps across different versions by taking the union of leaks detected in static and dynamic analysis and then comparing those leaks with the disclosure of PII in privacy policies.

The above studies have significantly contributed to the understanding of privacy leaks in Android apps by performing a *static* or *dynamic* analysis of apps. However, these studies have been conducted at a snapshot and have not been comprehensively evaluated over time. Given that the mobile ecosystem is continuously evolving, measurement studies longitudinally starting in the present may not comprehensively illuminate privacy policy violation trends to improve compliance. To fill this research gap, our study incorporates a longitudinal analysis spanning eight years. Furthermore, static analysis alone cannot deal with dynamically loaded code and obfuscated function calls, while dynamic analysis is prone to miss part of the function calls in the app. Our methodology raises the bar in the analysis of privacy leaks by combining static and dynamic analysis to capture most—if not all—of the leaks in Android apps. To our knowledge, the comprehensive longitudinal view on the Android app compliance with privacy policies has not been explored yet.

7 Concluding Remarks

In this paper, we analyzed the 5,240 historical versions, from 2008 to 2016, of 268 popular Android apps and investigated their compliance with the app privacy policies (P.P.). Our study found that most apps follow practices that contravene what is declared in their privacy policy. Our results also showed that apps are becoming more prone to violating their privacy policy than before, as we observed that the percentage of released app versions that comply with their privacy policy is steadily declining over time. In particular, the newer app versions disclose fewer PII collections in their P.P. and share more of a user’s private information through practices not mentioned in the P.P. This trend is of primary concern to users, especially considering that P.P. remains the cornerstone for protecting online privacy.

As future work, we aim to extend our study to recent years to evaluate the change in the data disclosure and collection (or sharing) practices caused by the GDPR. We also plan to study the similarity between the privacy policy of non-compliant apps to determine if third-party privacy policy generators have created their privacy policy. In this manner, app developers can be notified of the breach since app developers are often not experts in policy compliance and privacy laws.

References

1. “Complaint In the Matter of Snapchat, Inc.” <https://www.ftc.gov/system/files/documents/cases/141231snapchatcmpt.pdf>, 12 2014, [Accessed: 15-Oct-2021].
2. “Decision and Order In the Matter of Snapchat, Inc.” <https://www.ftc.gov/system/files/documents/cases/141231snapchatdo.pdf>, 12 2014, [Accessed: 15-Oct-2021].
3. “Video Social Networking App Musical.ly Agrees to Settle FTC Allegations That it Violated Children’s Privacy Law,” <https://www.ftc.gov/news-events/press-releases/2019/02/video-social-networking-app-musically-agrees-settle-ftc>, 2019, [Accessed: 15-Oct-2021].
4. “TikTok To Pay \$92 Million To Settle Class-Action Suit Over ‘Theft’ Of Personal Data,” <https://www.npr.org/2021/02/25/971460327>, 2021, [Accessed: 15-Oct-2021].
5. “TikTok Child Privacy Lawsuit Alleges Misuse of Personal Data of Millions of Minors,” <https://www.cpomagazine.com/data-protection/tiktok-child-privacy-lawsuit-alleges-misuse-of-personal-data-of-millions-of-minors>, 2021, [Accessed: 15-Oct-2021].
6. “1 in 5 children’s Google Play Apps breach Children’s Online Privacy Protection Act rules,” <https://www.comparitech.com/blog/vpn-privacy/app-coppa-study>, 2021, [Accessed: 15-Oct-2021].
7. “California Legislative Information - Internet Privacy Requirements,” https://leginfo.legislature.ca.gov/faces/codes_displaySection.html, 2003, [Accessed: 15-Oct-2021].
8. “California Legislative Information - AB-375 Privacy: personal information: businesses,” <https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml>, 2018, [Accessed: 15-Oct-2021].
9. “Children’s Online Privacy Protection Act of 1998 (COPPA),” <https://www.govtrack.us/congress/bills/105/hr4328>, 1998, [Accessed: 15-Oct-2021].
10. “EU Regulation - protection of natural persons with regard to the processing of personal data and on the free movement of such data,” <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, 2016, [Accessed: 15-Oct-2021].
11. “Data Protection Act 2018,” <https://www.legislation.gov.uk/ukpga/2018/12/contents>, 2018, [Accessed: 15-Oct-2021].
12. “The Privacy Act - OAIC,” <https://www.oaic.gov.au/privacy/the-privacy-act>, 1988, [Accessed: 15-Oct-2021].
13. C. Jensen, C. Potts, and C. Jensen, “Privacy practices of internet users: Self-reports versus observed behavior,” *International Journal of Human-Computer Studies*, vol. 63, no. 1-2, pp. 203–227, 2005.
14. A. M. McDonald and L. F. Cranor, “The cost of reading privacy policies,” *I/S: A Journal of Law and Policy for the Information Society (Isjlp)*, vol. 4, pp. 543–568, 2008.
15. C. Jensen and C. Potts, “Privacy Policies as Decision-Making Tools: An Evaluation of Online Privacy Notices.” New York, NY, USA: Association for Computing Machinery, 2004.
16. L. F. Cranor, “Necessary but not sufficient: Standardized mechanisms for privacy notice and choice,” *Journal on Telecommunications and High Technology Law*, vol. 10, pp. 273–307, 2012.

17. R. Slavin, X. Wang, M. B. Hosseini, J. Hester, R. Krishnan, J. Bhatia, T. D. Breaux, and J. Niu, "Toward a framework for detecting privacy policy violations in android application code." New York, NY, USA: Association for Computing Machinery, 2016.
18. X. Wang, X. Qin, M. B. Hosseini, R. Slavin, T. D. Breaux, and J. Niu, "Guileak: Tracing privacy policy claims on user input data for android applications." New York, NY, USA: Association for Computing Machinery, 2018.
19. I. Reyes, P. Wijesekera, J. Reardon, A. Elazari Bar On, A. Razaghpanah, N. Vallina-Rodriguez, and S. Egelman, "Won't Somebody Think of the Children? Examining COPPA Compliance at Scale," *Proceedings on Privacy Enhancing Technologies*, 2018.
20. E. Okoyomon, N. Samarin, P. Wijesekera, A. Elazari Bar On, N. Vallina-Rodriguez, I. Reyes, Á. Feal, and S. Egelman, "On the ridiculousness of notice and consent: Contradictions in app privacy policies," 2019.
21. C. Han, I. Reyes, Á. Feal, J. Reardon, P. Wijesekera, N. Vallina-Rodriguez, A. Elazar, K. A. Bamberger, and S. Egelman, "The price is (not) right: Comparing privacy in free and paid apps," *Proceedings on Privacy Enhancing Technologies*, vol. 2020, no. 3, 2020.
22. S. Zimmeck, P. Story, A. Ravichander, D. Smullen, Z. Wang, J. Reidenberg, N. C. Russell, and N. Sadeh, "MAPS: Scaling privacy compliance analysis to a million apps," in *19th Privacy Enhancing Technologies Symposium (PETS 2019)*, vol. 3. Stockholm, Sweden: Sciendo, July 2019, pp. 66–86.
23. B. Andow, S. Y. Mahmud, J. Whitaker, W. Enck, B. Reaves, K. Singh, and S. Egelman, "Actions Speak Louder than Words: Entity-Sensitive Privacy Policy and Data Flow Analysis with PoliCheck," in *USENIX Security Symposium*, 2020.
24. "Google Play Unofficial Python API," <https://github.com/NoMore201/googleplay-api>, [Accessed: 15-Oct-2021].
25. "User Data — Privacy, Security, and Deception - Developer Policy Center," <https://play.google.com/intl/en-US/about/privacy-security-deception/user-data>, [Accessed: 15-Oct-2021].
26. "Internet Archive: Digital Library of Free & Borrowable Books, Movies, Music & Wayback Machine," <http://archive.org>, [Accessed: 15-Oct-2021].
27. S. S. Hashmi, M. Ikram, and M. A. Kaafar, "A Longitudinal Analysis of Online Ad-Blocking Blacklists," in *Proceedings of the IEEE 44th LCN Symposium on Emerging Topics in Networking*, Osnabrück, Germany, 2019, pp. 158–165.
28. "Help with the EU user consent policy - Company - Google," <https://www.google.com/about/company/user-consent-policy-help>, 2018, [Accessed: 15-Oct-2021].
29. "Momento: Time Travel," <http://timetravel.mementoweb.org>, 2020, [Accessed: 15-Oct-2021].
30. S. Wilson, F. Schaub, A. A. Dara, F. Liu, S. Cherivirala, P. Giovanni Leon, M. Schaarup Andersen, S. Zimmeck, K. M. Sathyendra, N. C. Russell, T. B. Norton, E. Hovy, J. Reidenberg, and N. Sadeh, "The creation and analysis of a website privacy policy corpus," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016.
31. F. Liu, S. Wilson, P. Story, S. Zimmeck, and N. Sadeh, "Towards Automatic Classification of Privacy Policy Text," School of Computer Science Carnegie Mellon University, Pittsburgh, USA, Tech. Rep. CMU-ISR-17-118R, 06 2018.
32. "Tf-idf :: A Single-Page Tutorial - Information Retrieval and Text Mining," <http://www.tfidf.com>, 2020.

33. A. R. D. S. Z. W. J. R. N. C. R. Peter Story, Sebastian Zimmeck and N. Sadeh, "Natural Language Processing for Mobile App Privacy Compliance," in *AAAI Spring Symposium on Privacy Enhancing AI and Language Technologies (PAL)*, 03 2019.
34. J. Beel, B. Gipp, S. Langer, and C. Breitingner, "Research-paper recommender systems: a literature survey," *International Journal on Digital Libraries*, vol. 17, no. 4, pp. 305–338, 2016.
35. G. Chakraborty, M. Pagolu, and S. Garla, *Text Mining and Analysis: Practical Methods, Examples, and Case Studies Using SAS*. SAS Institute, 2014.
36. "Convert a collection of raw documents to a matrix of TF-IDF features," https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html, [Accessed: 15-Oct-2021].
37. "APP-350 Corpus," <https://usableprivacy.org/data>, 2019, [Accessed: 15-Oct-2021].
38. "One-vs-the-rest (OvR) multiclass/multilabel strategy," <https://scikit-learn.org/stable/modules/generated/sklearn.multiclass.OneVsRestClassifier.html>, [Accessed: 15-Oct-2021].
39. "Naive Bayes classifier for multinomial models," https://scikit-learn.org/stable/modules/generated/sklearn.naive_bayes.MultinomialNB.html, [Accessed: 15-Oct-2021].
40. "Logistic Regression (aka logit, MaxEnt) classifier," https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html, [Accessed: 15-Oct-2021].
41. "Linear Support Vector Classification (SVC)," <https://scikit-learn.org/stable/modules/generated/sklearn.svm.LinearSVC.html>, [Accessed: 15-Oct-2021].
42. "Apktool - A tool for reverse engineering Android apk files," <https://ibotpeaches.github.io/Apktool/>, [Accessed: 15-Oct-2021].
43. "Api reference—android developers," <https://developer.android.com/reference>, [Accessed: 15-Oct-2021].
44. "Java documentation: Naming a Package," <https://docs.oracle.com/javase/tutorial/java/package/namingpkgs.html>, [Accessed: 15-Oct-2021].
45. R. Binns, U. Lyngs, M. Van Kleek, J. Zhao, T. Libert, and N. Shadbolt, "Third party tracking in the mobile ecosystem." New York, NY, USA: Association for Computing Machinery, 2018.
46. M. Lindorfer, M. Neugschwandtner, L. Weichselbaum, Y. Fratantonio, V. van der Veen, and C. Platzer, "ANDRUBIS - 1,000,000 Apps Later: A View on Current Android Malware Behaviors," in *Proceedings - 3rd International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security, BADGERS 2014*, 2014, pp. 3–17.
47. J. Ren, M. Lindorfer, D. J. Dubois, A. Rao, D. R. Choffnes, and N. Vallina-Rodriguez, "Bug fixes, improvements, ... and privacy leaks: A longitudinal study of pii leaks across android app versions," in *Proceedings of Network and Distributed System Security Symposium*, 2018.
48. "GitHub - Eyasics/recon: Personal Information Exfiltration Detection Using Machine Learning," <https://github.com/Eyasics/recon>, 2016, [Accessed: 15-Oct-2021].
49. J. Ren, A. Rao, M. Lindorfer, A. Legout, and D. Choffnes, "ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic," in *Proceedings of the 14th*

- Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 361–374.
50. “ReCon - Controlled Experiments Code and Data,” <https://recon.meddle.mobi/codeanddata.html>, 2016, [Accessed: 15-Oct-2021].
 51. “ReCon - Format validation and String manipulation,” <https://github.com/Eyasics/recon/blob/master/code/src/meddle/RString.java>, 2016, [Accessed: 15-Oct-2021].
 52. “Android’s UI/Application Excerciser Monkey,” <https://developer.android.com/studio/test/monkey>, [Accessed: 15-Oct-2021].
 53. “mitmproxy - an interactive HTTPS proxy,” <https://mitmproxy.org>, [Accessed: 15-Oct-2021].
 54. “Understanding Dual Stacking of IPv4 and IPv6 Unicast Addresses,” https://www.juniper.net/documentation/en_US/junos/topics/concept/dual-stack-understanding.html, 2020.
 55. H. Zheng, D. Li, B. Liang, X. Zeng, W. Zheng, Y. Deng, W. Lam, W. Yang, and T. Xie, “Automated test input generation for android: Towards getting there in an industrial case.” IEEE Press, 2017.
 56. C. Gibler, J. Crussell, J. Erickson, and H. Chen, “Androidleaks: Automatically detecting potential privacy leaks in android applications on a large scale,” in *Trust and Trustworthy Computing*, S. Katzenbeisser, E. Weippl, L. J. Camp, M. Volkamer, M. Reiter, and X. Zhang, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 291–307.
 57. N. Viennot, E. Garcia, and J. Nieh, “A measurement study of google play.” New York, NY, USA: Association for Computing Machinery, 2014.
 58. “Fitbit home-page on Google Play,” <https://play.google.com/store/apps/details?id=com.fitbit.FitbitMobile>, [Accessed: 15-Oct-2021].
 59. S. S. Hashmi, M. Ikram, and S. Smith, “On optimization of ad-blocking lists for mobile devices,” in *Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2019, pp. 220–227.
 60. M. Ikram, P. Beaume, and M. A. Kâafar, “Dadidroid: An obfuscation resilient tool for detecting android malware via weighted directed call graph modelling,” in *Proceedings of the 16th International Joint Conference on e-Business and Telecommunications, ICETE 2019 - Volume 2: SECRYPT, Prague, Czech Republic, July 26-28, 2019*, M. S. Obaidat and P. Samarati, Eds. SciTePress, 2019, pp. 211–219.
 61. S.-K. Jo, M. Ikram, I. Jung, W. Ryu, and J. Kim, “Power efficient clustering for wireless multimedia sensor network,” *International Journal of Distributed Sensor Networks*, vol. 10, no. 4, p. 148595, 2014.
 62. M. Ikram and M. A. Kaafar, “A first look at mobile ad-blocking apps,” in *Network Computing and Applications (NCA), 2017 IEEE 16th International Symposium on*, 2017, pp. 1–8.
 63. B. Z. H. Zhao, M. Ikram, H. J. Asghar, M. A. Kaafar, A. Chaabane, and K. Thilakarathna, “A decade of mal-activity reporting: A retrospective analysis of internet malicious activity blacklists,” in *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*, 2019, pp. 193–205.
 64. T. Libert, “An automated approach to auditing disclosure of third-party data collection in website privacy policies.” Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, 2018.

8 Appendix

A Apps Selection

For the privacy policy analysis, we obtain 3,151 privacy policies and map them to 405 apps (comprising 7,998 APKs). Among the 3,151 privacy policies, 2,455 contain at least one *valid* segment and are mapped to 329 apps. From the 405 apps, we successfully download and analyze (static + dynamic) 350 apps comprising 7,741 APKs (or app versions) from the unofficial Google Play API. For each APK, we obtain its release date and assign a privacy policy based on the closest timestamp after the release date. For the compliance analysis, we only consider the APKs with at least one valid segment in their assigned privacy policy. Among the 7,741 APKs (spanning 350 apps) that we analyzed, 5,240 APKs (spanning 268 apps) satisfy the criterion and are considered for compliance analysis. It is possible for a given app to appear in privacy policy analysis and leaks analysis but not in compliance analysis. For example, an app has four versions (v1.1, v1.2, v1.3, and v1.4) and three unique privacy policies (pp1, pp2, and pp3). Suppose that pp1 is assigned to v1.1 and v1.2, and pp3 is assigned to v1.3 and v1.4. If pp2 contains valid segment(s), but pp1 and pp3 do not contain any valid segment, then the given app will not feature in the compliance analysis.

B Date difference

Figure 7 shows the time difference in months between the date when the privacy policy was crawled by the Wayback Machine and the release date of the app. 42.6% APKs have a time difference of fewer than 12 months with their mapped privacy policy date.




Fig. 7. Difference in months between the privacy policy date and APK release date.

C Label Descriptions and Classifiers Performance

Table 6 describes the various PII labels used in our privacy policy classification. Table 7 shows the performance of various machine learning classification algorithms on the test data-set of APP-350 corpus [37].

Table 6. Descriptions of PII disclosures in the privacy policies that are collected/shared.

Data Type (PII)	Description
Contact	Unspecified contact data.
Contact_Address_Book	Contact data from a user’s address book on the phone.
Contact_City	User’s city.
Contact_E-Mail_Address	User’s e-mail.
Contact_Password	User’s password.
Contact_Phone_Number	User’s phone number.
Contact_Postal_Address	User’s postal address.
Contact_ZIP	User’s ZIP code.
Demographic	User’s unspecified demographic data.
Demographic_Age	User’s age (including birth date and age range).
Demographic_Gender	User’s gender.
Identifier	User’s unspecified identifiers.
Identifier_Ad_ID	User’s ad ID (such as the Google Ad ID).
Identifier_Cookie	User’s HTTP cookies, flash cookies, pixel tags, or similar identifiers.
Identifier_Device_ID	User’s device ID (such as the Android ID).
Identifier_IMEI	User’s IMEI (International Mobile Equipment Identity).
Identifier_IMSI	User’s IMSI (International Mobile Subscriber Identity).
Identifier_IP_Address	User’s IP address.
Identifier_MAC	User’s MAC address.
Identifier_Mobile_Carrier	User’s mobile carrier name or other mobile carrier identifier.
Identifier_SIM_Serial	User’s SIM serial number.
Identifier_SSID_BSSID	User’s SSID or BSSID.
Location	User’s unspecified location data.
Location_Bluetooth	User’s Bluetooth location data.
Location_Cell_Tower	User’s cell tower location data.
Location_GPS	User’s GPS location data.
Location_IP_Address	User’s IP location data.
Location_WiFi	User’s WiFi location data.

Table 7. Performance of different classifiers. Multinomial Naive Bayes (MNB), Logistic Regression (LReg), and Linear Support Vector Classifier (SVC). The accuracy scores for SVC range from 91.16% to 100%.

Classifiers	accuracy (%)			precision (%)		
	MNB	LReg	SVC	MNB	LReg	SVC
Contact	98.6	98.7	98.8	0	77.8	65.4
Contact_Address_Book	98.8	99.3	99.4	0	79.5	80.5
Contact_City	99.6	99.8	99.8	0	81.8	73.3
Contact_E-Mail_Address	95.9	97.3	97.2	84.8	84.8	83
Contact_Password	98.8	99.2	99.4	0	70.7	78.6
Contact_Phone_Number	97	98.8	98.8	9.1	86.7	83.6
Contact_Postal_Address	98	98.6	99	0	68.5	80.3
Contact_ZIP	99.4	99.8	99.8	0	94.1	94.4
Demographic	98.7	99.6	99.6	25	79.2	78.9
Demographic_Age	98.3	99.3	99.4	16.7	89.3	90
Demographic_Gender	98.9	99.6	99.6	16.7	79.3	79.3
Identifier	99.1	99.1	99.1	0	0	66.7
Identifier_Ad_ID	98.8	99.7	99.8	0	97.7	97.8
Identifier_Cookie	95.4	98.7	98.5	90.3	88	88.6
Identifier_Device_ID	97.9	99	99	73.6	86.7	83.7
Identifier_IMEI	99.4	99.8	99.9	50	92.6	93.1
Identifier_IMSI	99.9	99.9	99.9	0	100	100
Identifier_IP_Address	99.1	99.4	99.4	90.1	94.1	94.8
Identifier_MAC	99.3	99.9	99.9	0	92.6	96.3
Identifier_Mobile_Carrier	99.4	99.7	99.7	0	100	85.7
Identifier_SIM_Serial	99.8	99.8	99.8	0	44.4	56.3
Identifier_SSID_BSSID	99.8	99.8	99.9	0	0	100
Location	96.7	98.8	98.8	80.9	86.5	87
Location_Bluetooth	99.4	99.8	99.8	0	80	81
Location_Cell_Tower	99.5	99.9	99.9	25	79.2	81.8
Location_GPS	99	99.7	99.7	64.7	90.2	88.7
Location_IP_Address	99.3	99.6	99.6	0	85.7	78.3
Location_WiFi	99.6	99.8	99.8	70	77.8	83.3
Performed	90.5	91.7	91.2	91.7	90.9	86.3
Not_Performed	98.6	98.9	98.7	3.2	34.2	30.5
1stParty	91	92.5	92.3	81	85	82.7
3rdParty	94.9	96.7	96.7	53.3	73.8	71.6
average	98.06	98.82	98.82	28.94	77.22	81.92